

Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market

Sara Fisher Ellison
M.I.T. and CESifo

Christopher M. Snyder
Dartmouth College and NBER

Hongkai Zhang
Pandora Media Inc.

May 2018

Abstract: We study price dynamics for computer components sold on a price-comparison website. Our fine-grained data—a year of hourly price data for scores of rival retailers—allow us to estimate a dynamic model of competition, backing out structural estimates of managerial frictions. The estimated frictions are substantial, concentrated in the act of monitoring market conditions rather than entering a new price. We use our model to simulate the counterfactual gains from automated price setting and other managerial changes. Coupled with supporting reduced-form statistical evidence, our analysis provides a window into the process of managerial price setting and the microfoundation of pricing inertia, issues of growing interest in industrial organization and macroeconomics.

Keywords: price dynamics, managerial costs, behavioral economics, structural estimation

JEL Codes: L11 (Production, Pricing, and Market Structure); C73 (Stochastic and Dynamic Games), D22 (Firm Behavior: Empirical Analysis), L81 (e-Commerce)

Contact Information: Ellison: Department of Economics, M.I.T., 50 Memorial Drive, Cambridge, MA 02142; email sellison@mit.edu. Snyder: Department of Economics, Dartmouth College, 301 Rockefeller Hall, Hanover, NH 03755; email chris.snyder@dartmouth.edu. Zhang: Pandora Media Inc., 2100 Franklin Street, Oakland CA 94612; email hongkaicol@gmail.com.

Acknowledgments: The authors are grateful to Michael Baye, Steve Berry, Judy Chevalier, Oliver Compte, Glenn Ellison, Jakub Kastl, Peter Klenow, John Leahy, Emi Nakamura, Whitney Newey, Mo Xiao, and Jidong Zhou for their insightful advice; to Masao Fukui for outstanding research assistance; and to seminar participants at Bologna, Cornell, Harvard, M.I.T., Paris School of Economics, Toronto, and Yale and conference participants at the ASSA Meetings, HSE–Perm International Conference on Applied Research in Economics, International Industrial Organization Conference (Boston), International Symposium on Recent Developments in Econometric Theory with Applications in Honor of Jerry Hausman (Xiamen University), NBER Summer Institute (IT and Digitization and Economic Fluctuations Meetings), and Tuck School of Business Winter Industrial Organization Workshop for helpful comments.

1. Introduction

Managerial frictions can be broadly construed as a cost associated with the manager's overlooked but essential input into the operation of a firm. Textbook industrial-organization models leave little room for managerial frictions. Consider, for example, the ubiquitous pricing decision. Firms in most industries have some discretion over their output price, which may require dynamic adjustment to changing supply and demand conditions. When does the manager make these price changes? By how much? Textbook industrial-organization models would say price is changed precisely when and by the amount dictated by strategic considerations based on all available information. Other economic subdisciplines have paid more attention to the frictions in this pricing process. Macroeconomists have developed a class of models in which price inertia plays a key role in the monetary transmission mechanism, leading to a large theoretical and empirical macro literature studying a range of frictions such as menu costs, the costs of monitoring rivals' prices, managerial inattention, and the implicit contracts that firms have with their customers to maintain prices. Organizational economists have long been interested in the internal processes that firms use to overcome the complexities associated with price setting in practice, dating back to Simon's (1962) claim that "Price setting involves an enormous burden of information gathering and computation that precludes the use of any but simple rules of thumb as guiding principles," and to Cyert and March's (1963) classic book documenting the rule of thumb used by a department store to price its merchandise.

In this paper, we help bridge the gap between these literatures by extending methods used in industrial organization to estimate dynamic models to provide some of the first structural estimates of the sort of managerial frictions posited by macro and organizational economists. We study the pricing decisions of a set of rival firms selling computer components in an online marketplace, Pricewatch, which ranks the lowest-price firms in the most prominent positions, generating the highest sales. Price changes were relatively frequent in this market (the average spell between price changes lasting about a week) yet far from continuous. Initial reduced-form analysis coupled with information from manager interviews provide evidence that managerial frictions played a role in this price inertia, motivating us to propose a structural model to quantify these frictions. Our fine-grained data (hourly price observations for scores of firms over a year) combined with fluctuating market conditions and the jockeying for price rank generate a rich sample of price changes, an ideal environment to estimate a dynamic model of competitive price adjustment. We use the model to back out structural estimates of managerial frictions, separately identifying a manager's cost of monitoring the market from the menu cost he or she faces when entering a new price. These structural estimates provide insight into the microfoundations of the observed price inertia, a perennial concern of macro and organizational economists, but also of interest to industrial-organization economists concerned with price dynamics.

We estimate our dynamic, structural model following the two-step method of Bajari, Benkard, and Levin (2007) (BBL) for estimating structural parameters in the profit function. The first step involves estimating reduced-form strategies constituting a Markov-perfect equilibrium. The second step involves finding structural parameters in the profit function that rationalize the estimated strategies such that no player can improve

his payoff by deviating to some other strategy. To allow for feasible estimation in our complex setting and to accommodate the possibility of boundedly rational managers—a natural possibility in our setting in which managerial frictions are the focus—we make several modifications to BBL’s procedure. The reduced-form policy functions estimated in the first step are not required to constitute a Markov-perfect equilibrium based on the universe of state variables but rather are taken to be rule-of-thumb strategies based on a subset of salient ones. In the second stage, we search for structural parameters rationalizing the estimated rule of thumb as undominated by any deviation in the class of admissible rules of thumb. Rather than requiring that the firm’s strategy to be optimal in every state, we impose the weaker requirement that the firm chooses a rule of thumb *ex ante* maximizing its expected present discounted value of the flow of surplus over the play of the game, where the expectation is taken over the distribution of salient states experienced by the firm. The computational burden that the full state space would impose is daunting. Just considering permutations of the scores of sample firms across price ranks, putting aside many other payoff-relevant variables (costs, margins, distances between firms in price space, etc.), there are billions times more such permutations in a single period than grains of sand on earth. Even if estimation were feasible for the econometrician, it may be unrealistic to assume that managers could compute the best responses each instant conditional on such a complex state space, especially in our setting in which the simple act of monitoring rivals’ prices will present a substantial managerial friction, to say nothing of performing wildly complicated calculations. Our approach—as one of the first attempts to accommodate boundedly rational players in a dynamic structural model—presents its own challenges, to be successful requiring the specification of a parsimonious policy functions providing a good fit to observed play. We will thus need to pay considerable attention to the specification of the functional form and covariates included in the policy function and to the measurement of the policy function’s goodness of fit. To improve the fit of the policy functions, we modify the treatment of firm heterogeneity in BBL’s framework. Instead of estimating the mean, variance, and other parameters characterizing the continuous distribution of structural parameters across firms, we postulate discrete types of firms, use machine-learning techniques to partition firms into the types *ex ante*, and estimate separate structural parameters for each type. The key benefit of assuming several discrete types is that we can improve the flexibility of the policy-function specification by allowing its coefficients to vary freely across types. For parsimony, we end up postulating three types; if we instead try the structural estimation with a single policy function, its fit and the structural estimates are quite poor.

Consistent with recent work microfounding price stickiness, including contributions to the macro theory literature by Alvarez, Lippi, and Paciello (2011) and Bonomo *et al.* (2015), we build two separate managerial frictions into the structural model, (a) a cost of monitoring market conditions to determine the appropriate timing and size of a price change and (b) a menu cost, i.e., the physical cost of recording a new price. Though we do not observe monitoring, we are still able to identify the cost of this latent action through the exclusion restriction that certain payoff-relevant variables are only available to the manager after monitoring. We find substantial monitoring costs—around \$68 per episode for the most prominent firm type—but virtually no menu costs. The division of managerial frictions matters for firm profit. In the counterfactual exercise in

which these costs are reversed, with costless monitoring and a \$68 menu cost, the firm’s net profit rises substantially because the firm is able to tailor its price policy more closely to market conditions. In another counterfactual exercise, we document substantial profit gains from an increasingly automated price-setting policy (which firms went on to implement after our sample period). Our structural estimates confirm the findings of Alvarez, Lippi, and Paciello (2015), who use survey data to calibrate their previously cited dual-cost model, finding that the monitoring cost is three times the menu cost of changing price.

A limit of our approach is that our conclusions are derived from a small and specialized market, not a broader set of markets representative of the U.S. economy. But even then, we would argue that more general lessons can be drawn from our findings here, and the detail of the results more than outweighs the disadvantage of specificity. For instance, measuring price stickiness in an online marketplace is a useful complement to much of the existing empirical evidence, which so far has concentrated on brick-and-mortar stores.¹ Also, we would argue that simply understanding the price-setting mechanisms in one market very well, as we can with our unusually detailed data, can help us predict which factors are likely to be important in many other markets. Finally, structural estimation would be difficult without the fine-grained, market-specific data provided by our setting. Our structural approach enables us to simulate counterfactual managerial costs for a firm, including higher costs that the manager of a brick-and-mortar store might face.

The paper is structured as follows. Section 2 reviews related literatures. Section 3 provides background on the Pricewatch market. Section 4 describes the data. Section 5 highlights several important reduced-form results from the data. Besides motivating the setup of the structural model, these reduced-form results contribute independent insights into price-changing behavior, which have the appeal of being non-technical and fairly assumption-free. The rest of the paper is devoted to specification and estimation of the structural model. Section 6 specifies the three components of our structural model: the value, profit, and policy functions. Section 7 presents estimates of the policy function and assesses the goodness of its fit. Section 8 discusses identification of the structural parameters, presents the structural estimates, and provides a sensitivity analysis. Section 9 conducts counterfactual simulations showing how a firm’s pricing behavior changes when managerial costs are varied from the estimated levels. The last section concludes. An appendix provides technical details on the structural estimation.

2. Literature Review

In this section, we provide more detail on where our paper fits in the literatures of several subfields. Methodologically, our paper lies in the industrial-organization literature, in particular the literature that structurally estimates dynamic games, including seminal papers by Aguirregabiria and Mira (2007); Pakes, Ostrovsky, and Berry (2007); Pakes, *et al.* (2015); and BBL cited in the introduction. Our estimation strategy is based on BBL, modified to allow boundedly rational managers to choose a rule-of-thumb policy. Our paper studies the same Pricewatch market and uses the same data as Ellison and Ellison (2009a, 2009b). Those papers

¹Exceptions include Arbatskaya and Baye (2004), Chakrabarti and Scholnick (2005), Lønneman and Wintr (2006), Gorodnichenko and Weber (2015), and the “Billion Prices Project” (Cavallo and Rigobon 2012) in progress.

focus on the demand side, estimating demand elasticities with respect to price, rank, avoided sales tax, and geographic proximity between buyer and seller. We shift the focus to the supply side, analyzing pricing dynamics, allowing for the strategic interaction among firms, and structurally estimating managerial frictions. We will take the demand estimates from Ellison and Ellison (2009a) as an input into our analysis.

Our paper touches on several other strands of the industrial-organization literature. The field has a distinguished history of interest in price rigidity, dating at least back to Gardiner Means's testimony to Congress testimony about inflexibility of prices during the Depression (Means 1935, documented in Terkel 1970). Stigler and Kindahl (1970) and Carlton (1986) bypass the Consumer Price Index (CPI) computed by the Bureau of Labor Statistics (BLS) and indices computed by other administrative agencies, drilling into the underlying transactions data to establish a body of stylized facts about price rigidity.

Price rigidity is but one issue under the broader heading of dynamic pricing. Another issue that has received attention among empirical industrial-organization economists is documenting the occurrence of Edgeworth cycles (Edgeworth 1925, formalized by Maskin and Tirole 1988), in which firms gradually undercut each other until they reach the zero-profit level where they stay until one relents, resulting in a dramatic price rise. Edelman and Ostrovsky (2010) and Zhang (2010) document Edgeworth cycles in sponsored-search and online-advertising auctions, respectively.² Our setting resembles these Internet auctions in that firms continuously "bid" for favorable rank positions by cutting their margins—indeed winning the rank-1 position often requires the firm to have a negative margin of price over cost. Our reduced-form analysis will document cycles in ranks, but in reverse of the Edgeworth pattern: a given firm's rank gradually rises as others undercut as they adjust their prices for secular declines in cost, punctuated by sharp drops in the given firm's rank when the manager attends to pricing, readjusting to a target rank.

A large industrial-organization literature documents price dispersion in homogeneous, consumer-good markets ranging from retail gasoline (Barron, Taylor, and Umbeck 2004; Hosken, McMillan, and Taylor 2008; Lewis 2008) to books (Clay, *et al.* 2002) to general retail (Lach 2002). Closest to our setting is Baye, Morgan, and Scholten (2004), who document price dispersion in an price-comparison website using rich data collected with a web-scraping technology. Although we also document price dispersion in our reduced-form analysis, this is not one of our core results, which are structural estimates of managerial frictions. However, our core results can be viewed as contributing to the price-dispersion literature. Our finding that substantial managerial frictions contribute to price stickiness provides a novel explanation for equilibrium price dispersion. Most explanations rely on costly price search à la Stahl (1989), but managerial frictions could also generate a distribution of prices for homogeneous goods in equilibrium. That is, our structural estimates raise the possibility that supply-side rather than demand-side frictions drive price dispersion.

Our paper is related to several subfields outside of industrial organization. Substantively, our analysis is related to a large literature in macroeconomics on sticky prices, including notable papers by Blinder *et*

²A number of papers attempt to characterize Edgeworth cycles and other interesting pricing dynamics in gasoline markets, including Castanias and Johnson (1993); Eckert and West (2004); Noel (2007a, 2007b, 2008); Hosken, McMillan, and Taylor (2008); Atkinson (2009); Lewis (2009); Wang (2009); and Doyle, Muehlegger, and Samphantharak (2010).

al. (1998); Bils and Klenow (2004); Nakamura and Steinsson (2008); and Eichenbaum, Jaimovich, and Rebelo (2011). As summarized in Klenow and Malin’s (2011) handbook chapter, this literature uses a range of different data sources from manager surveys to supermarket scanner data to price data used by the BLS to compute the CPI. These are typically reduced-form studies documenting facts about frequency and size of price changes as well as identifying sectors with the most price rigidity (nondurable goods, processed goods, goods with noncyclical demand, and goods sold in more concentrated markets). Comparing our methodology to this literature’s, we also document descriptive facts in our from an initial reduced-form analysis, but our main goal is structural estimation of managerial frictions, which this literature does not undertake. Comparing our data to this literature’s, each has certain advantages. The macro studies tend to use comprehensive datasets of retail prices, representative of the economy, whereas we look at a single product. The tradeoff is that their datasets have less detail for each product. For example, BLS price data is collected monthly, so price movements within a month are not recorded. We have hourly data, and we do, in fact, observe many price changes within a month. We observe prices for most of the rivals operating in the market, allowing us to estimate firms’ reactions to their rivals’ price movements. The BLS does not attempt to sample all the substitutes in a product space. The nature of our market and the length of our time series results in our recording dozens of price changes for each firm, allowing us to estimate a detailed policy function for price changes at the firm level. We also have unusually detailed auxiliary data on wholesale costs and quantities, integral to the structural estimation.³

Ours is closest to a suite of papers from macro as well as empirical trade that back out managerial frictions in price setting by, in effect, measuring the imperfect pass through of input-price shocks (generated by exchange-rate or commodity-price movements, depending on the study) to product prices. Two pioneering studies estimate models of price setting by a single agent with costly adjustment, Slade (1998) in the market for crackers sold at grocery stores, Davis and Hamilton (2004) in the market for wholesale gasoline. More recently, Nakamura and Zerom (2010) and Goldberg and Hellerstein (2013) move beyond single-agent models to oligopolies, using a static oligopoly model including menu costs. Nakamura and Zerom (2010) find a menu cost of \$7,000 in their calibration exercise, far greater than any managerial friction we estimate, though their figure is still a small percentage (less than 0.25%) of revenue. Goldberg and Hellerstein (2013) estimate a menu cost of \$60 to \$230, or 1% to 3% of revenue in their structural estimation of the static oligopoly model, in the range of our estimated monitoring cost, considerably higher than the negligible menu cost we found in our online setting. Similar to these previous papers, we use variation in input prices, in our case the wholesale price of memory chips, to provide crucial exogenous variation behind

³A recent addition to the macro literature, a working paper by Gorodnichenko and Weber (2015), studies pricing in an online platform serving many retailers, generating data of similar richness as ours. Like most of this literature, the authors do not estimate parameters governing firm interactions or managerial decisions, so, although their empirical setting is similar to ours, their goals and methods are quite different. A study of the market for used cars in Germany in Artinger and Gigerenzer’s (2012) working paper spans several literatures, macro as well as organizational economics. In the spirit of the macro literature, in particular Blinder *et al.* (1998), the first part of their paper presents descriptive evidence from extensive manager surveys. The second part of their paper gathers online price data from hundreds of car dealers to test Simon’s (1955) model of aspirational pricing. While we share their goal of understanding what drives firms to change online prices, the direction of our analysis is quite different, estimating a model of price changing by firms and structural parameters capturing managerial frictions.

our estimation of managerial frictions. We diverge from these seminal papers in two ways, by incorporating the dynamics of firm interactions and by distinguishing two types of managerial costs, a crucial distinction, given that we find menu costs to be dominated by a different form of managerial friction, monitoring costs.

3. Empirical Setting

Our empirical setting is the online marketplace for computer components mediated by Pricewatch, previously studied by Ellison and Ellison (2009a, 2009b); see those papers for additional details on the empirical setting. During the 2000–01 period during which our data were collected, the Pricewatch marketplace was composed of a large number of small, undifferentiated e-retailers selling memory upgrades, CPUs, and other computer parts. These retailers tended to run bare-boned operations with spartan offices, little or no advertising, rudimentary websites, and no venture capital. A large fraction of their customers came through Pricewatch. Instead of click-through fees, retailers paid Pricewatch a monthly fee to list products. Customers could use Pricewatch to locate a product in one of two ways, either typing a product description into a search box or running through a multi-layered menu to select one of a number of predefined product categories. For example, clicking on “System Memory” and then on “PC100 128MB SDRAM DIMM” would return a list of products in that category sorted from cheapest to most expensive in a format with twelve listings per page. The full list for a pre-defined category could span dozens of pages with listings from scores of retailers. Figure 1 contains the first page of a typical list for the memory module in our study, downloaded during our sample period.

The Pricewatch ranking exhibited substantial churn from day to day and even from hour to hour. Figure 2 illustrates the movement in prices (top row of panels) and ranks (bottom row of panels) for three representative retailers of PC100 128MB memory modules in our sample during a representative month. The first and third retailers changed price six times during the month, the second retailer twice. Even the second firm’s slower rate of price change is quite rapid compared to findings in the macro literature from U.S. Bureau of Labor Statistics data that the median length of spells between price changes of 4.3 months (Bils and Klenow, 2004). While frequent, retailers’ price changes were still far from continuous. When firms were not changing prices, their ranks continued to move, bumped up or down by rival price changes, leading a firm’s rank to fluctuate much more than its price.

Based on information from a detailed interview of a manager of one of the retailers participating on Pricewatch, we can identify several possible reasons for this churn in the rankings. First, wholesale prices could be quite volatile for some products. Retailers would receive wholesale-price quotes via daily emails, and they often fluctuated from day to day. Short-term fluctuations could be up or down, but as these are electronic components, the long-term trend was downward. The daily price quotes were relevant to the manager’s operations as they typically carried little or no inventory, ordering enough to cover just the sales since the previous day’s order. A second reason for turnover in the rankings was that managers did not continuously monitor each individual product’s rank on the Pricewatch website, making the instantaneous

changes needed to maintain a particular rank. Retailers typically offered scores of products in different categories on Pricewatch; it would be impractical for a manager to continuously monitor all of them even he or she attended to no other managerial tasks. Certain high-volume products, such as the specific memory modules we study, could merit more attention, but this might involve checking at most one or two times a day. During our study period, managers had to enter price changes manually into Pricewatch’s database, as automated price setting was not introduced until 2002. Each adjustment would thus involve a fixed cost in both determining and entering the appropriate price.

A retailer’s rank on Pricewatch was a key determinant of its sales and profits. The first panel of Figure 3, derived from demand estimates from Ellison and Ellison (2009a), shows how a retailer’s daily sales vary with rank. The bulk of sales go to the two or three lowest-priced retailers in this market, but positive sales still accrue to many additional retailers on the list. For later reference in our calculation of firm profits, we will label the curve $Q(Rank_{it})$. A significant source of profit in this market came from the “upselling” strategy documented by Ellison and Ellison (2009a), by which a firm attracts potential customers with low prices for the “base” memory module, but then tries to induce them to upgrade to a more expensive one. The second panel of Figure 3 again uses results from Ellison and Ellison (2009a) to back out an estimate of the hourly profit from this upselling strategy as a function of the firm’s rank. The shape results from two opposing forces: at higher ranks the firm has fewer potential customers, but the customers it does attract are advantageously selected be more likely to upgrade. For later reference in our calculation of profits, which will incorporate returns from the upselling strategy, we will label the curve $U(Rank_{it})$.

Returning to Figure 2, it illustrates another feature of the Pricewatch market that will play a large role in our analysis: heterogeneity in retailers’ pricing strategies. We have already seen that retailers varied in the frequency of price change. They also appear to have targeted different segments of the ranking, with the first firm maintaining low ranks, the second allowing itself to drift into middle ranks, and the third content with high ranks. To accommodate this heterogeneity in the later estimation, we will allow the parameters to vary freely across different types of firms. In fact, looking ahead to Section 5.1, the three firms in Figure 2 are representatives of the three types into which we will ultimately classify our sample using machine-learning methods. This analytical treatment allows us to be agnostic about the source of the firms’ strategy heterogeneity, whether differences in firm costs, such as the cost incurred by a manager to monitor market conditions or to compute and enter a price change, or differences in a firm’s ability to convert customers attracted by its rank position into sales of the base product or upgraded product (i.e., differences in the $Q(Rank_{it})$ or $U(Rank_{it})$ function).

To summarize the main take-aways from this brief overview of the market, competition among retailers for Pricewatch rank, a key competitive variable, led to frequent although far from continuous price changes. Our interview information suggested that this pricing inertia was due in part to managerial frictions, motivating further reduced-form work providing clearer documentation of managerial inertia and motivating structural estimation of a model of managerial costs involving in monitoring the market and changing price. Visual evidence of substantial heterogeneity in firms’ pricing strategies will motivate our accommodating

this heterogeneity in the estimation in a flexible way.

4. Data

Our data come from Ellison and Ellison (2009a). The authors scraped information on the first two pages of listings (12 listings per page for a total of the 24 lowest-price listings) from the Pricewatch website for several computer components. We focus on the category of 128MB PC100 memory modules because it is the most active and highest volume of the categories collected and because of the homogeneity of the products in this category. The authors recorded the name of the firm (as well as other information about the firm such as its location), the name of the specific product, and its price. By linking names of firms and products over time, we are able to trace pricing strategies of individual firms for individual products (taking the conservative approach of assuming that a change in the product’s name indicates a change in product offering). The authors scraped this continuously updated information every hour from May 2000 to May 2001 (with a few interruptions).

Ellison and Ellison (2009a, 2009b) supplemented the Pricewatch data with proprietary data from a retailer who sold through Pricewatch. This retailer provided information on its quantity sold and wholesale acquisition cost. We take this cost to be common across retailers justified by the fact that the typical retailer carried little inventory, did not have long-term contracts with suppliers, and had access to a similar set of wholesalers.

A large number of firms made brief appearances on the Pricewatch lists. Since we are interested in the dynamics of firms’ pricing patterns, we study firms that were present for at least 1,000 hours during the year (approximately one-eighth of our sample period) and changed price while staying on the list at least once. For a small number of firms who had multiple products on the first two pages of Pricewatch simultaneously during some periods, we excluded their observations during those periods. We were left with 43 firms appearing at some point during the year, at most 24 present on the first two pages of Pricewatch at any particular moment. Although the excluded retailers do not constitute observations, we do use them to compute relevant state variables for rivals including rank, density of neighboring firms, etc.

Based on these data, we created a number of variables to describe factors that might be important to firms’ decisions about timing and magnitude of price changes. Figure 2 demonstrated the importance of the rank of that firm on the Pricewatch list for firm outcomes. We also included margin, length of time since its last price change, number of times a firm has been “bumped” (i.e., had its rank changed involuntarily) since its last price change, and so forth. Table 1 provides a description of these variables and summary statistics.

Most of the variables can be understood from the definitions in the table, but a few require additional explanation. *Placement* measures where a firm is between the next lower- and next higher-priced firms in price space. For example, if three consecutive firms were charging \$85, \$86, and \$88, the value for *Placement* for the middle firm would be $0.33 = (86 - 85) / (88 - 85)$. *Density* is a measure of the crowding of firms in the price space around a particular firm. It is defined as the difference between the price of the

next higher-priced firm minus the price of the firm three spaces below divided by 4. For example, if five consecutive firms charged \$84, \$84, \$85, \$86, and \$88, the value for *Density* for the firm charging \$86 would be $1 = (88 - 84)/4$. *QuantityBump* reflects relative changes in a retailer’s order flow caused by being bumped from its rank. It is calculated using the $Q(Rank_{it})$ function in Figure 2, in particular proportional to $\ln(Q(Rank_{it})/Q(Rank_{it'}))$, where t is the current period and t' is the period in which the retailer last changed price. *CostTrend* and *CostVol* are computed by regressing the previous two weeks of costs on a time trend and using the estimated coefficient as a measure of the trend and the square root of the estimated error variance as a measure of the volatility. The definitions of the remaining variables are self-explanatory.

Turning to the descriptive statistics, we see that the average price for a memory module in our sample was \$69, with a considerable standard deviation of 35.1. Most of this variation is over time, with prices typically above \$100 at the beginning of the period and down in the \$20s by the end, mirroring a large decline in the wholesale cost of these modules. The mean spell between price changes was 117.55 hours—about five days—similar to what we saw for the first and third of the representative firms in Figure 2. Wholesale cost showed a strong downward trend over our sample period, falling an average \$0.19 per day, but was quite volatile.

5. Reduced-Form Evidence

Before diving into the structural model, we pause in this section to report several results from a reduced-form analysis of the data. Besides motivating the setup of the structural model, the results contribute independent insights into price-changing behavior, which have the appeal of being non-technical and fairly assumption-free.

5.1. Firm Heterogeneity

To accommodate the strategy heterogeneity illustrated by the representative firms in Figure 2 in our subsequent empirical analysis, we will adopt a compromise between combining all firms in one sample and performing separate firm-by-firm analysis. Separate firm-by-firm analysis sacrifices power, perhaps unnecessarily; results in a profusion of parameters that are hard to digest; and, most importantly, selects a non-random set of firms, eliminating almost all of the less active firms due to too few observations. To balance these three concerns with a desire to allow heterogeneity in our model, we will classify firms into a small number of strategic types and allowed the model parameters to differ freely across the types.

To partition the firms into types, we employed a popular machine-learning technique, cluster analysis.⁴ The first step is to select a set of dimensions along which the firms could be differentiated. We chose seven variables that we judged were instruments under the firms’ control rather than outcomes depending on

⁴See Romesburg (2004) for a textbook treatment. To be clear, the cluster analysis referred to here is not the same as “clustering the standard errors,” which is the familiar way of adjusting standard errors for correlation among related observations (which as indicated in the table notes we do throughout, clustering by firm).

external factors.⁵ Next, the variables are standardized so that each has a standard deviation of 1, preventing the variable with the largest variance from dominating the assignment. Starting with every firm its own cluster, the algorithm proceeds by identifying which clusters are most similar, measured by the sum of Euclidean distances between all firms in the two clusters, and iteratively combines them.⁶ We iterated until three clusters were left, which we thought achieved a balance between spanning most of the important firm heterogeneity and still obtaining enough of each type enough to analyze empirically with some confidence. We will allow estimates of most structural parameters to freely vary across the three clusters, henceforth referred to as types $\tau = 1, 2, 3$.

Table 2 provides variable means by each of the three firm types. The 22 firms of type 1 generally occupy the lower price ranks and change prices relatively frequently. The eight firms of type 2 occupy the middle price ranks and tend to change prices infrequently. The 13 firms of type 3 generally charge the highest prices but are more active in changing prices than type 2. This characterization echoes what we saw in Figure 2 for the three firms representing each type. The means for *Margin* show that type 1 firms earn the lowest margins (at least captured by this measure), followed by type 2, followed by type 3. The means of *Rank* show the same pattern, with type 1 occupying the lowest ranks, followed by type 2, followed by type 3. The means of *SinceChange* show that type 1 and type 3 change price more than twice as often as type 2.

5.2. Distribution of Price Changes

Figure 4 provides more refined evidence on the distribution of the size and spells of price changes by firm type. Panel A provides histograms of the size of price changes. The bar for zero change has been omitted for readability since firms did not change price in the vast majority of hours. Note that the height of the other bars are unconditional masses, i.e., not conditioned on a price change. The distributions are roughly unimodal for the three types, the mass falling off for larger magnitude price changes. The mode for all types is at a price reduction of \$1. While there is mass for price increases as well as reductions, more mass is on the reduction side, consistent with the downward trend in wholesale costs throughout our sample. A characteristic feature of type 2 firms is that they change price only infrequently, which shows up in the histograms as relatively little overall mass in the histogram compared to the other types. The ratio of tail mass to mode mass is higher for type 2 firms than others, suggesting that their infrequent price changes tend to be larger in magnitude than other types’.

Panel A can be related to the theoretical and empirical literatures from macroeconomics on price stickiness. Our histograms are roughly unimodal, inconsistent with theoretical models such as Golosov and Lucas

⁵The variables include the firm’s target *Rank* and *Placement*, the firm’s mean values of *NumBump*, *SinceChange*, and *FirstPage*, and the firm’s variance of *SinceChange*. We also included the fraction of time the firm was present in our sample. Firm *i*’s target value of a given variable is the mean of that variable computed for the subset of periods that immediately follow a price change by *i*.

⁶The method of starting with each item in a separate cluster and combining them until the target is reached is called the agglomeration method. The use of Euclidean distance (the sum of squared differences in the standardized variables) to measure difference and the criterion of combining clusters that have the smallest sum of squared differences is called Ward’s method. These are the standard options in Stata 14’s cluster command, which we used to perform the cluster analysis.

(2007) that, under the assumption of cheap monitoring and expensive price changing, predict bimodal distributions, with most of the mass on large price changes in the tails than on small price changes in the center of the distribution. We will actually provide a microfoundation for the unimodal price distributions shown here when we estimate managerial costs. In particular, we estimate a high cost of monitoring and low cost of price changing. The empirical macro literature has found unimodal distributions in other settings such as Figure 2 in Midrigan (2011) and Figure II in Klenow and Kryvtsov (2008). The magnitudes of our firms' price changes are smaller than in Klenow and Kryvtsov (2008), who report over 35% of price changes exceeding 10% of the initial price in absolute value, while only 7% of our price changes are this large.

Panel B plots a kernel density estimate of the distribution of spells between price changes. The distribution looks similar for types 1 and 3. These frequent price changers necessarily have shorter spells. Most of the mass is concentrated in spells shorter than 200 hours. The mass for type 2 firms spreads out more uniformly over longer spells, even as high as 600 hours.

The distribution of spells that we report has a shape similar to those previously documented, but with a compressed scale. For instance, the spell distributions for our type 1 and type 3 firms look quite similar to Figure IV in Klenow and Kryvtsov (2008) and Figure VIII in Nakamura and Steinsson (2008), but instead of a time scale stretching to 12 or 18 months respectively, ours covers just one month. The difference could be due in part to the nature of the data: much of the previous literature used monthly price data, missing more frequent price changes that our hourly data would pick up. However, the analysis of survey data in Blinder *et al.* (1998) finds that a large majority of managers change the prices of their most important products no more than twice in a year. The managers of our firms are simply more active than those in these other studies, perhaps due to volatile wholesale costs or low managerial costs.

Panel C plots the absolute value of the price change as a nonparametric function of spell using Cleveland's (1979) locally weighted regression smoothing (LOWESS). The graph for type 1 firms shows an upward slope, with larger changes made after longer spells. The type 2 graph is also initially upward sloped. There is a downturn for longer spells, but this non-monotonicity is estimated from too few observations to be conclusive. The type 3 graph shows a flat relationship between size and spell. The graph for type 2 is higher than the other types', indicating that the magnitude of type 2's price changes tend to be larger even conditional on spell. Although there is not a single pattern emerging here that can be used to make a clean comparison, our results at least are not at odds with what the previous literature has found for the relationship between spell length and size of price change, such as Figure VIII in Klenow and Kryvtsov (2008).

5.3. Managerial Inertia

We have seen that price changes, while frequent, were far from continuous, on the order of once per week rather than each hour. This pricing inertia could be due in principle to many factors. The integer constraint on prices could combine with only slowly moving market forces to produce the infrequent price changes. Perhaps firms are reluctant to engender a rival response to a price cut. Figure 5 provides evidence that the inertia is due at least in part to costs of managerial activity.

The figure plots the residual probability of a price change (after partialling out other covariates as controls) during each hour measured in Eastern Time, estimated separately for retailers operating on the East and West Coasts. Retailers supply a national market via Pricewatch, so if there were no managerial costs, presumably the timing of activity would be similar on the two coasts, responding to the same national demand factors. In fact the probability functions peak at different points, at 11 a.m. for East Coast and 8 p.m. for West Coast retailers.

Interestingly, the peak for the West Coast retailers is not simply shifted by the three-hour difference between Eastern Time and the local time for West Coast retailers, as might be expected if the retailers on the two coasts served separate markets but faced the same pattern of managerial costs during the day. Instead, the peak on the West Coast is shifted by ten hours, to 8 p.m. Eastern Time, which is 5 p.m. in their local time. One explanation, supported by the interview subject, is that the market has already been operating for a few hours by the time the West Coast managers arrive at work, so they are well-advised to set prices in the evening before they leave. The evening might also be a less busy time for them since orders might have started falling off at least from customers in the East. In contrast, the East Coast manager would arrive to a very slow order flow at 8 a.m. Eastern Time and have the leisure to adjust prices at that point before orders picked up for the day. In the absence of managerial costs, it is difficult to rationalize why retailers on the two coasts would pick the opposite ends of the day to do most of their price changes.

6. Structural Model

This section presents the model that will be the basis for our structural estimation. We begin in the next subsection by discussing the key object in dynamic structural estimation, the firm's value function. The subsections after that discuss the components that go into computation of the value function.

6.1. Value Function

Firm i participates in the market each period from the current one, t , until it exits at time T_i . Its objective function is the present discounted value of the stream of profits given by the value function

$$V_i(s_t; \sigma, \theta) = E \left[\sum_{k=t}^{T_i} \delta^k \pi_{ik} \right], \quad (1)$$

where s_t represents the state of the game in period t , σ is the vector over i of firms' strategies σ_i , δ is the discount factor,

$$\pi_{it} = \pi(s_t, \sigma(s_t), \theta) \quad (2)$$

is the static profit function, reflecting the payoff from one period (one hour in our empirical setting) of play, and θ is a vector of parameters. The expectation is taken over the distribution of all possible game play and evolution of private shocks starting from s_t .

The focus of this study is on obtaining structural estimates of θ , which will include measures of managerial costs, upselling profits, and other variables of central interest. In essence, we will compare firm i 's value function when it plays equilibrium strategy σ_i to that when it plays some deviation $\tilde{\sigma}_i$, maintaining rivals' play as specified by σ . The estimated θ will be those values minimizing violations of the dominance of σ_i over $\tilde{\sigma}_i$. As in BBL, we are not required to solve for the equilibrium. Instead, we can observe the equilibrium from the data. Still, we need to compute the value function in and out of equilibrium, which in turn requires three components: (a) specification of the profit function π_{it} ; (b) an estimate of firm i 's policy function, $\hat{\sigma}_i(s_t)$, which, following BBL, we will use in place of the equilibrium strategy $\sigma_i(s_t)$; and (c) treatment of the expectations operator. The remainder of this section will be devoted to specifying the profit function and policy function. Expectations will be computed by averaging the present discounted value of the profit stream from many simulated runs of the market, as described in more detail in Section 8 on the structural estimation.

For tractability and to better suit our empirical setting, our model will depart in several ways from the standard BBL framework. We touch on the departures here but discuss them in more detail below in the relevant sections. First, either because of limited information, cognition, or memory, the manager is assumed to only be able to consider a restricted set $\hat{S}$ of the overall state space. For example, rather than keeping track of the whole price history for each rival, manager i may view his current rank as an adequate summary statistic of the combined effect of these actions. Second, the manager is assumed to follow a rule-of-thumb policy that is a function of the restricted state space $\hat{S}$. He does optimize, not over the price each moment, but over the long-run choice of this policy, maximizing the present discounted value of profits over the set of admissible policies PF . Given the overwhelming complexity of the period-by-period optimization problem involving the full state space, it is unlikely that managers were fully neoclassical—all the more in our setting in which we have qualitatively documented substantial managerial frictions, which we aim to quantify. Our approach has the advantage of being able to accommodate behavioral managers. The disadvantage of model is that it will be misspecified if managers are more neoclassical or behavioral in a different way than we allow. It will thus be essential for us to specify a class of policy functions that can closely match observed firm behavior. We will devote careful attention to this specification and to gauging its resulting goodness of fit.

6.2. Profit Function

Our detailed specification of the profit function π_{it} draws on our rich information about the business strategies of the firms operating in this market and the costs they face, as well as institutional details and estimates from Ellison and Ellison (2009a):

$$\pi_{it} = Base_{it} + Upsell_{it} - \mu_{\tau} Monitor_{it} - \chi_{\tau} Change_{it}, \quad (3)$$

where

$$Base_{it} = Q(Rank_{it})(Price_{it} - Cost_t) \quad (4)$$

$$Upsell_{it} = U(Rank_{it}). \quad (5)$$

The first term $Base_{it}$ accounts for the profits from the sale of the base version of the memory modules, equaling the quantity sold times the margin per unit. The quantity sold is the function $Q(Rank_{it})$ from Figure 3A. The second term $Upsell_{it}$ accounts for the upselling strategy discussed in Section 3, whereby the firm attracts potential customers with low prices for the base product but then induces some of them to upgrade to more expensive versions of the memory module. It is given by the function $U(Rank_{it})$ from Figure 3B.⁷

The final two terms in (3) reflect the two types of managerial costs that we seek to estimate, both of which vary by firm type τ . The coefficient μ_τ on the indicator $Monitor_{it}$ for whether firm i monitors in period t will provide an estimate of the cost of monitoring. The last term involves $Change_{it}$, which is our label for the indicator $1\{\Delta_{it} \neq 0\}$ for whether firm i changes price in period t , where we define $\Delta_{it} = Price_{it} - Price_{i,t-1}$. The coefficient χ_τ on this indicator will provide an estimate of the cost of price change over and above any monitoring cost. We do not observe monitoring, but, just as we will generate simulated values for $Rank_{it}$, $Price_{it}$, and other variables when we compute expectations of the value function, we can simulate $Monitor_{it}$ using the model of monitoring behavior embedded in the policy function discussed in the next subsection.

The specification of managerial costs μ_τ and χ_τ as fixed parameters represents a departure from BBL's framework, which specifies a distribution over structural parameters with a mean and variance to be estimated. Our approach of estimating a single fixed parameter eases the computational burden, important given we allow the parameters to differ across the three types τ . We also have several conceptual reasons for our approach. The policy function specified in the next subsection already incorporates arrival of opportunities for managerial actions based on both stochastic as well as deterministic factors. Our estimates of the structural parameters can be thought of as the marginal cost of taking the relevant managerial action in opportune periods dictated by the policy function. Implicit in this formulation is the presumption that managerial actions are not taken in other periods in part because of high, perhaps prohibitive, draws for random managerial costs. To explain the relative rarity of price changes observed on a hourly basis, an estimated distribution of managerial costs would have to be skewed toward large values compared to our estimate for opportune periods. Estimating the distribution of managerial costs would require specifying the functional form for the distribution and the hour-to-hour correlation structure, which could be quite complex. Our policy function captures the time-series properties of opportunities for managerial action simply, by including covariates such as *SinceChange*, *Night*, and *Weekend*, among others.

⁷We normalize the profit function for a firm that moves off of the first two Pricewatch pages to \$4.70, our projection of a firm's hourly profit at rank 25 based on Ellison and Ellison's (2009a) parametric estimates for ranks 1–24. The structural estimates are robust to the choice of this normalization.

6.3. Policy Function

The next piece needed for structural estimation is the policy functions, or models of firms' strategies. An estimate of this model provides the $\hat{\sigma}_i(s_t)$ that will be substituted for equilibrium strategies $\sigma_i(s_t)$ to compute the value function (1) in the simulations. The model reflects the reality that price changes in this market are infrequent relative to our data frequency. To a first order, the best prediction of the firm's price next hour is its current price. We will thus focus on modeling the timing and size of price change episodes, with the firm maintaining the price outside of these episodes. Consistent with our belief that firms are not engaging in complicated calculations of optimal policy based on hundreds of state variables each hour, we want the model to be simple, streamlined, and reflective of just the variables that firms are likely to be able to monitor and process. Also, we want the model to be a predictive empirical description of what firms actually do.

We thus model price changes as coming from a two-step process. The manager knows some components of the market state vector at all times, information he receives essentially "for free." Based on these state variables, the manager's first step is to decide whether to attend to the market to gather information and perform calculations needed for a pricing decision. We will call this behavior "monitoring" and denote the decision to do so with the indicator function $Monitor_{it}$. In our setting, we think of monitoring as computing quantities like percentage markups and visiting the Pricewatch website to gather the relevant information, involving an opportunity cost of cognition and time. Through monitoring, the manager gains additional information on state variables, including current rank and the distribution of competitors' prices, and computes the new desired price. If the new desired price is different from the current price and the costs of changing it are justified, then he enters it in the Pricewatch form, and changes it on his own website, again involving costs in terms of cognition and time. We will call this second step "price changing."

The two-step process can account for periods of excess inertia, during which the manager keeps price constant even though market conditions would warrant a price change. Inertia can come from three sources. First, the manager may not be aware of the changed market conditions because he did not monitor. Second, the benefit from making a desired price change, especially a small price change, may not justify the managerial cost of entering it. Third, if the desired price change is smaller than a whole dollar unit in which Pricewatch prices are denominated, price may stay constant. The two-step process is also consistent with anecdotal evidence from our interview subject and broader survey evidence (see Blinder et al., 1998) that managers often monitor market conditions including rival prices without changing their own price.

We specify the manager's latent desire to monitor, $Monitor_{it}^*$, as

$$Monitor_{it}^* = X_{it}\alpha_\tau + e_{it}, \quad (6)$$

where X_{it} is a vector of explanatory variables, α_τ is a vector of coefficients to be estimated, which are allowed to differ across firm types τ , and e_{it} is an error term. If $Monitor_{it}^* \geq 0$, then the firm monitors; i.e., $Monitor_{it} = 1$. Otherwise, if $Monitor_{it}^* < 0$, then the firm does not monitor; i.e., $Monitor_{it} = 0$.

In essence, equation (6) embodies the manager's forecast of the costs and benefits of monitoring, so,

performance, the explanatory variables can only include state variables known by the manager before monitoring. In addition, the variables must be important shifters of either the cost or benefit of monitoring. We specify the following parsimonious list:

$$X_{it} = \left(Night_t, Weekend_t, CostVol_t, CostTrend_t^+, |CostTrend_t^-|, QuantityBump_{it}^+, |QuantityBump_{it}^-|, \ln SinceChange_{it}, (\ln SinceChange_{it})^2 \right). \quad (7)$$

We assume the manager is automatically aware of the time and day. The variables $Night_t$ and $Weekend_t$ are included to reflect that the cost of monitoring varies in a predictable way over a week: the cost of monitoring at 2am on Sunday morning might be high if that's when a manager typically sleeps, and the benefit might be low because few sales would be made around that time anyhow. The manager is also assumed to be aware of the day's wholesale cost—recall that he receives emails every day from the wholesalers—and can glean volatility and trends from the pattern of costs over the past couple of weeks. Presumably the gains to monitoring are greater the more conditions including costs are fluctuating. We include $CostVol_t$ to capture the magnitude of recent unpredictable fluctuations and $CostTrend_t$ to capture recent trends. Both rapidly rising and rapidly falling costs would lead the manager to monitor more. To allow asymmetry⁸ in the concern for rising or falling costs, a rising cost trend, $CostTrend_t^+ = CostTrend_t \times 1\{CostTrend_t > 0\}$, enters (6) separately from a falling cost trend, $CostTrend_t^- = |CostTrend_t| \times 1\{CostTrend_t < 0\}$. The latter variable appears in absolute value so that we would anticipate the two variables' coefficients to have the same sign if not magnitude.

We assume that the manager is roughly aware of changes in his order flow resulting from being bumped in the ranks and will be more likely to monitor if there has been a large change, whether an increase or decrease. Thus (6) includes $QuantityBump_{it}$. Recall that this variable is computed by translating the current rank and rank at the previous price change into a quantity change using the function in Figure 3A. This predicted change in order flow is a proxy for the quantity signal the manager observes. The proxy diverges from the signal because the firm's actual sales depend on random market fluctuations on top of any predictable effect of a rank change. The proxy also diverges from the signal because the manager may only be vaguely aware of actual sales in any given hour. Again, to allow for asymmetries, increases in order flow, $QuantityBump_{it}^+$, enter separately from decreases, $QuantityBump_{it}^-$.

The last set of variables, functions of $SinceChange_{it}$, enter in a flexible, nonlinear way, allowing for various patterns of managerial attention, including monitoring the market at regular intervals as well as periods of intense monitoring, in which several price changes may follow in succession, followed by periods of inattention during which price stays constant independent of market conditions. Including this variable can help us tease out time dependence from state dependence in price monitoring behavior.

⁸Some evidence suggests that prices could be stickier in one direction versus the other. For instance, Borenstein, Cameron, and Gilbert (1997) identify an asymmetry in the response of wholesale gasoline prices to cost increases versus decreases.

Conditional on monitoring, the manager may decide to change price based on the information acquired. Let $\Delta_{it}^* = Price_{it}^* - Price_{i,t-1}$ denote the size of the price change the manager would desire if price were a continuous variable and the menu cost χ_τ of changing price were ignored just for period t . Assume this latent variable is given by

$$\Delta_{it}^* = \begin{cases} Z_{it}\beta_\tau + u_{it} & \text{if } Monitor_{it} = 1 \\ 0 & \text{if } Monitor_{it} = 0, \end{cases} \quad (8)$$

where Z_{it} is a vector of explanatory variables, β_τ is a vector of coefficients to be estimated, which again are allowed to differ across firm types τ , and u_{it} is an error term. The actual price change $\Delta_{it} = Price_{it} - Price_{i,t-1}$ may diverge from the latent price change Δ_{it}^* for two reasons. First, rather than being continuous, prices were denominated in whole dollars on Pricewatch. Second, because changing price is not in fact costless, the manager must weigh the benefits of changing price if cost were no object reflected by Δ_{it}^* against the menu cost χ_τ . Assuming the manager's willingness to pay to change price is an increasing function of the size of the desired price change, we can relate the observed price change Δ_{it} to the actual Δ_{it}^* by specifying cut points along the real line

$$\dots < C_\tau^{-5} < C_\tau^{-4} < C_\tau^{-3} < C_\tau^{-2} < C_\tau^{-1} < C_\tau^1 < C_\tau^2 < C_\tau^3 < C_\tau^4 < C_\tau^5 < \dots \quad (9)$$

Then

$$\Delta_{it} = \begin{cases} k & \text{if } \Delta_{it}^* \in (C_\tau^k, C_\tau^{k+1}), \\ -k & \text{if } \Delta_{it}^* \in (C_\tau^{-(k+1)}, C_\tau^{-k}). \end{cases} \quad (10)$$

Thus, for example, an observed price increase of \$1 corresponds to a latent price change satisfying $\Delta_{it}^* \in (C_\tau^1, C_\tau^2)$. If the manager monitored but did not change price, then $\Delta_{it}^* \in (C_\tau^{-1}, C_\tau^1)$. To deal with the fact that many elements of the manager's decision to change price—the function mapping the size of his desired price change into his willingness to pay to change price, the distribution of the error term u_{it} , of course the menu cost χ_τ , which is yet to be estimated—we adopt a flexible specification for the C_τ^k , allowing them to be free parameters estimated in a similar way as the cut points in an ordered probit, and allowing them to differ across firm types τ .

The explanatory variables can include state variables the manager learns as a result of monitoring in addition to those known before. We specify the following parsimonious set:

$$Z_{it} = \left(CostTrend_t, CostChange_{it}, Margin_{it}, NumBump_{it}, \right. \\ \left. Density_{it} \times NumBump_{it}, Placement_{it}, Rank_{it}, RankOne_{it} \right). \quad (11)$$

The higher are forecasted costs, $CostTrend_t$, and the more costs have risen since firm i last changed its price, $CostChange_{it}$, the higher the firm's desired price. $Margin_{it}$ may also factor into price changes, a firm with

low margins being more likely to increase price and one whose margins are already high less likely to further increase price and more likely to lower.

The remaining variables in Z_{it} are the sort of state variables revealed by monitoring. After visiting the Pricewatch website, the firm learns its current rank and thus the number of ranks it was bumped since the last price change. The firm can use this information to return itself to its desired rank, so decreasing price if it was bumped up in the ranks (gauged by a positive $NumBump_{it}$) and increasing price if it was bumped down. A low-ranked firm may have less of an incentive to cut price to increase its sales; this effect may be particularly strong for a firm occupying rank 1: further price reductions will only result in a small increase in sales because most of the demand elasticity is with respect to rank, which the firm cannot improve beyond 1. We, therefore, include an indicator, $RankOne_{it}$. We include $Density_{it}$ because the presence of a thicket of close competitors may affect pricing incentives. For example, in a dense price space, firms may have less incentive to increase price because this will result in a more severe rank change. Specifically, we include $Density_{it}$ interacted with $NumBump_{it}$ because density will likely matter more for firms that have a need to change price as proxied by a bump from the previous preferred ranking. A firm's $Placement_{it}$ between its nearest rivals will also affect its desired price in potentially complex ways.

The list of explanatory variables in Z_{it} deliberately excludes some of the variables that appear in the monitoring equation. For example, $Night_t$ is included in the monitoring equation to reflect the fact that checking the Pricewatch website at 2 a.m. would typically be more costly than during the workday. However $Night_t$ should have little effect on the desired price change conditional on monitoring because, conditional on already being on the website, changing price is no more difficult at 2 a.m. than 10 a.m. The same logic applies to $Weekend_t$. While the time since price was last changed may affect the desire to monitor—one possibility is that with more time for market conditions to change, more information is to be gained—conditional on the information gained through monitoring such as $Rank_{it}$ and $NumBump_{it}$, $SinceChange_{it}$ has no obvious role in price setting and thus is excluded from Z_{it} .

While it is easy to justify the inclusion of the explanatory variables in the respective functions, it may be harder to assert that this strategy model is sufficient to describe firms' behavior. Firms could have executed much more complex strategies, including nonlinear functions of the included variables, interactions of them, and additional variables such as the characteristics of neighboring firms in the Pricewatch ranking. One might worry whether a parametric policy function could adequately capture the intricacies involved in full-blown expected-value maximization each period, perhaps calling for more flexible methods such as random forests, which can systematically process a range of possible covariates and specifications, nesting our two-stage model.

We argue that our approach still has a number of appealing features in our empirical setting. First, our approach is computationally much less burdensome than a fully rational model. Second, it is reasonable to suppose managers used fairly simple rules of thumb to make the high-frequency decisions to monitor and change price rather than continually evaluating expected value functions based on scores of state variables each period. Such calculations would overwhelm the direct cost of checking websites and calculating and

typing in new prices involved in the managerial actions. Third, as discussed in Section 7.4, we have encouraging results from a goodness of fit test to determine the match between simulations from our estimated strategies and observations in the data. Finally, flexible machine-learning methods such as random forests are not panaceas. Although excellent at prediction, random forests often lack the useful economic interpretability of a parametric model. That said, incorporating machine-learning methods in dynamic structural estimation is a promising area for future research.

7. Estimation of the Policy Function

7.1. Estimation Details

Our model for the policy function embodied in equations (6)–(11) is known in the econometrics literature as a zero-inflated ordered probit (ZIOP). The term “zero-inflated” refers to the fact that there are more zeros—in our setting periods without a price change—than would be expected under a standard ordered probit. The overabundance of zeros in our setting results from the manager’s monitoring less than continuously because of the cost of monitoring. Thus the monitoring stage determines the degree of zero-inflation. Harris and Zhao (2007) demonstrate in Monte Carlo experiments that the maximum likelihood estimator of the ZIOP performs well in finite samples. While we adopted the ZIOP specification and chose the variables appearing in the each stage to best reflect our understanding of the empirical setting informed by manager interviews, formal Vuong (1989) tests of non-nested models reported in Section 7.4 support the inclusion of a monitoring stage to the price-change equation and the support our choice of variables to include in the different stages.

We estimate equations (6)–(11) jointly using maximum likelihood taking the errors e_{it} and u_{it} to be independent standard normal random variables.

Without the monitoring stage, the price-change stage of the model would simply be an ordered probit, where various intervals would correspond to various discrete price changes. With the monitoring stage, the likelihood of observations in which firm i does not change price at time t is

$$L(\Delta_{it} = 0) = \underbrace{1 - \Phi(X_{it}\alpha_\tau)}_{\Pr(\text{Monitor}_{it}=0)} + \underbrace{\Phi(X_{it}\alpha_\tau)}_{\Pr(\text{Monitor}_{it}=1)} \underbrace{\left[\Phi(C_\tau^1 - Z_{it}\beta_\tau) - \Phi(C_\tau^{-1} - Z_{it}\beta_\tau) \right]}_{\Pr(\Delta_{it}^* \in (C_\tau^{-1}, C_\tau^1))}, \quad (12)$$

where Φ is the standard normal distribution function. The likelihood of, for example, a k dollar price increase is

$$L(\Delta_{it} = k) = \underbrace{\Phi(X_{it}\alpha_\tau)}_{\Pr(\text{Monitor}_{it}=1)} \underbrace{\left[\Phi(C_\tau^{k+1} - Z_{it}\beta_\tau) - \Phi(C_\tau^k - Z_{it}\beta_\tau) \right]}_{\Pr(\Delta_{it}^* \in (C_\tau^k, C_\tau^{k+1}))}. \quad (13)$$

Adding the monitoring stage scales up the probability of no price change and scales down the probability of any given sized price change.

We group the few price reductions of \$5 or more together and similarly group the few price increases of

\$5 or more so that we only need to estimate thresholds down to C_τ^{-5} and up to C_τ^5 .

7.2. Identification

Though the dependent variable $Monitor_{it}$ in the monitoring equation (6) is not itself observable, the coefficients α_τ in this equation can still be estimated because $Monitor_{it}$ affects the distribution of observed price changes (or lack thereof) through the latent price change equation (8).

The monitoring and price-changing equations combine to produce a single price change, but the parameters in the separate stages can still be separately identified. Intuitively, parameters in the price-changing equation are identified from periods with active price changes. Fixing the covariates, a unique parameter vector will best match the distribution of price increases and decreases. The number of zeros may not be well matched; the mismatch identifies the parameters in the monitoring equation. For example, an excess of zeros is attributed to a reduced likelihood of monitoring since, had the firm had been monitoring, these opportunities to change price would not have been missed.

In Monte Carlo exercises, Harris and Zhao (2007) find that the ZIOP model is well identified even without exclusion restrictions, but exclusion restrictions can help avoid convergence problems and improve the precision of parameter estimates. Exclusion restrictions play a more central role in applications such as ours in which heteroskedasticity or non-normal errors cannot be ruled out, precluded by data limitations from estimating the error distribution nonparametrically. The problem is similar to that arising with the Heckman (1979) selection model, which is well known to perform poorly in the presence of non-normal or heteroskedastic errors without exclusion of some variables appearing in the selection equation from the main equation.

The natural exclusion restrictions we employ bolster identification. We include $Night_t$ and $Weekend_t$ in the monitoring equation but exclude them from the price-change equation because they relate more to the inconvenience of activity at those times than the optimal size of a price change. We also include $SinceChange_{it}$ in the monitoring equation, capturing the manager's guess of how far out of line the unadjusted price has become over time, but excluded from the price-change equation because other variables are sufficient statistics for the optimal price change once the manager is able to observe them through monitoring. These excluded variables will help identify true excess zeros (when these instrumental variables are high, say) from a distributional mode that is higher than expected due to non-normal or heteroskedastic errors.

7.3. Results

Table 3 presents estimates of α_τ , β_τ , and C_τ^k by firm type τ . Consider the results for type 1 firms first, as these represent the majority of firms, comprising the biggest sample. *Night* and *Weekend* show up as important determinants of whether a firm monitors. This result is not surprising but is a telling indicator of the importance of managerial attention. Managers evidently take evenings and weekends off and do not bother monitoring the market then. Managers are also more likely to monitor if wholesale cost is volatile, if there has either been a sharp trend in wholesale cost or a sharp change in order flow recently, or if there

has been a long spell since the last price change. These estimates are consistent with intuition and are quite precisely estimated for type 1 firms. Conditional on monitoring, a price increase for a type 1 firm is associated with a wholesale cost increase, a firm being bumped down, rank being low, rank being 1, and there being few firms close by in price space. Again, these results are consistent with intuition as well as our conversations with one of the firm managers.

A likelihood-ratio test rejects homogeneity of the policy function across types at better than the 1% level. While there are quantitative differences among the results for the firm types, there are generally not large qualitative differences, especially for the α_τ coefficients. Our framework is flexible enough to accommodate the difference in activeness between type 2 firms and the rest. In the monitoring equation, the coefficients for $\ln SinceChange$ and $(\ln SinceChange)^2$ of type 2 firms constitute a quadratic function of $\ln SinceChange$ that is decreasing when $SinceChange < 586$. So a type 2 firm would only monitor when it has accumulated enough change in *QuantityBump* or it has been a very long time since the last price change. Conditional on monitoring, a type 2 firm requires a more significant latent size of price change to trigger actual price changes because the size of the interval of no price change, $C_\tau^1 - C_\tau^{-1}$, is larger for $\tau = 2$ than other types.

7.4. Goodness of Fit

This subsection assesses the goodness of fit of our estimated policy function $\hat{\sigma}$, in particular whether it fits well enough to be suitable for the later structural estimation. As discussed in the next section, after substituting the specified profit function as well as estimates of the structural parameters, firm i 's value function $\hat{V}_i(\hat{\sigma}, \theta)$ reduces to a linear function of just a few aggregates, in particular, the total hours firm i spends at each rank and the total times it changes its price. Hence the goodness of fit of the estimated policy function can be assessed by comparing aggregates simulated from the policy function to their values in the actual data. We first provide an informal, visual assessment using a series of figures.

To provide some technical details on how the simulations were run, we start with a given firm type and consider all the observations at the it level for that type. For each it , we construct 20 simulated forward histories lasting 720 hours. The simulations use the actual market data for state variables where possible (e.g., for cost histories), simulating firm behavior by substituting current state variables as well as random draws for error terms e_{it} and u_{it} into equations (6) and (8) of the policy function. Averaging the aggregate over all the simulations for all it observations of that firm type produces a value that can be compared to the average in the actual data of 720-hour forward histories starting from each candidate observation. To be precise, the aggregates are discounted rather than simple sums, using the same annual discount factor of 0.95 used in the value functions in the structural estimation. Over the 720-hour horizon, discounting has a negligible effect on the results.

The first panel of Figure 6 compares simulated time spent at each rank (solid black curves) to actual (dashed grey curves) for the three firm types. The actual curves have quite different levels and shapes across firm types, yet the simulated curves are able to fit each quite well. Panel B compares the simulated number of price changes from the estimated policy function (the black squares) to the averages in the actual

data (the grey circles). Again, the fit is quite close, with the markers essentially overlapping and moving together across the types of firm: moderate for type 1, low for type 2, and high for type 3. The close fit is not surprising given the policy function is a reduced form that was estimated in part to maximize the likelihood of the observed frequency of price changes. However, the close fit was not guaranteed. The maximum likelihood estimation targeted size as well as number of price changes; we see that the joint estimation does not harm the fit for number alone. More importantly, the policy functions were estimated to fit individual behavior; letting their interaction on the market play out over a length of time could generate feedback causing behavior to diverge from actual outcomes. We see in the first panel that this is not the case. Taken together, the results from Figure 6 suggest that the estimated policy function will provide good estimates of the sums that are the essential inputs into the value functions in the structural estimation.

Figure 7 provides a more refined assessment of goodness of fit. The policy function should not only be able to fit the forward history on average across states but also fit the forward history in any state s in which the firm finds itself. The figure compares simulated to actual profit conditional on various dimensions of the state including various initial ranks in the first panel and initial margins in the second panel. (To save space, we just show the results for type 1 firms; the fit for types 2 and 3 is similar.) We focus on profit in this figure as opposed to time spent at each rank in the previous figure because profit is a convenient summary statistic for the distribution of times spent at each rank, allowing us to reduce the dimensionality of the graph. While we do not yet have all the components of profit π_{it} from equation (3)—the managerial-cost components are of course yet to be structurally estimated—we can estimate the first two components $Base_{it}$ and $Upsell_{it}$ using equations (4) and (5). Because monetary profits are not observed even in the actual data, they have to be estimated in both cases—using simulated prices in the former case and actual prices in the latter. The figure reports results binned first by rank and then by margin. Each point on the graph represents an average taken over it observations whose rank or margin at the point of observation (serving as the initial state s from which we simulate 720 hours forward) puts it into the relevant bin.

Panel A compares the simulated to actual monetary profit going forward for 720 hours conditional on rank in the initial period. The solid black curve for profit based on simulated prices closely matches the dashed grey one for profit based on actual prices over the whole range of the horizontal axis. Of course the closeness of the graphs could be a symptom, not of good fit, but of the stability of the environment, with initial rank correlating highly with profits at least over an horizon as short as 720 hours. To investigate this possibility, we have added a curve (lighter with dot markers) representing the naïve forecast that the firm earns the same profit in each of the 720 hours as it does in the first. This naïve forecast ends up overestimating profit conditional on top ranks, because it has not taken into account the other firms' reaction to firm i 's top rank, which is to undercut firm i . For similar reasons it ends up underestimating profit conditional on bottom ranks. Our estimated policy function fits dramatically better for both high and low ranks, suggesting that our policy function likely captures this dynamic correctly.

Panel B similarly compares simulated, actual, and naïve estimates of monetary profit conditional on a type 1 firm's margin (price minus wholesale cost in levels) in the initial state s . Again our estimated policy

simulation fares well, while the naïve prediction underestimates profit conditional on negative margins and overestimates profit conditional on positive margins because does not properly incorporate the firm’s future price adjustments and the cascade of rival responses. This comparison suggest that our policy function has likely captured this dynamic correctly.

Table 4 provides an additional assessment of the policy function’s fit using formal Vuong (1989) tests of model closeness. The Vuong (1989) test statistic Z compares the model’s likelihoods after subtracting a penalty for extra parameters. We compare our preferred model, the zero-inflated ordered probit (ZIOP) embodied in equations (6)–(11), labeled model A, to a series of alternatives, labeled B–D.

For model B, we start with our preferred specification but shift some variables from X_{it} in the monitoring equation to Z_{it} in the price-change equation; in particular we shift all the variables except $Night_t$ and $Weekend_t$ from X_{it} to Z_{it} . This comparison will test whether the careful choice of which equation the variables appear in matters or whether it just matters the variables are included somewhere in the specification. For model C, instead of shifting the variables from X_{it} to Z_{it} , we just omit them from the model entirely. This comparison will test whether the monitoring equation is mainly identified off of the timing variables $Night_t$ and $Weekend_t$ or whether other included variables contribute as well. Model D keeps the price-change equation as specified in the preferred Model A but omits monitoring equation (6). Thus, unlike models A–C, which are all ZIOP models, model D is a simple ordered probit (OP) model.

In cases of nested or partially overlapping models, the Vuong Z statistic has a complex distribution, related to a weighted sum of chi squares where the weights are eigenvalues of a matrix of conditional expectations, leading authors to provide a subjective evaluation of the level of the Z statistic rather than directly computing its p -value. We see that the test statistic is enormous for the comparison of the OP model omitting monitoring (model D) against any of the ZIOP models (models A–C). This suggests that including the monitoring equation produces a substantially better fit. The preferred ZIOP specification (model A) has a considerable Z statistic of 3.20 and 3.92 when tested against the alternatives that either shift (model B) or delete variables (model C) from the monitoring equation. This provides suggestive evidence that careful placement of the variables in the correct equations matters for fit.

8. Estimation of Structural Parameters

8.1. Identification

For tractability and to better suit our empirical setting, our model will depart in several ways from the standard BBL framework. We touch on the departures here but discuss them in more detail below in the relevant sections. First, either because of limited information, cognition, or memory, the manager is assumed to only be able to consider a restricted set $\hat{S}$ of the overall state space. For example, rather than keeping track of the whole price history for each rival, manager i may view his current rank as an adequate summary statistic of the combined effect of these actions. Second, the manager is assumed to follow a rule-of-thumb policy that is a function of the restricted state space $\hat{S}$. He does optimize, not over the price each moment,

but over the long-run choice of this policy, maximizing the present discounted value of profits over the set of admissible policies PF . Given the overwhelming complexity of the period-by-period optimization problem involving the full state space, it is unlikely that managers were fully neoclassical—all the more in our setting in which we have qualitatively documented substantial managerial frictions, which we aim to quantify. Our approach has the advantage of being able to accommodate behavioral managers. The disadvantage of model is that it will be misspecified if managers are more neoclassical or behavioral in a different way than we allow. This motivated our careful specification of the policy function to matching observed firm behavior and our assessment of its goodness of fit.

We follow the broad outlines of the BBL approach to estimate our structural parameters with a few modifications. While BBL assume the econometrician observes the game play of a Markov perfect Nash equilibrium, we assume a weaker solution concept such that the observed game play is a Nash equilibrium in which each player chooses a rule-of-thumb policy function based on a limited set of state variables to maximize the present discounted value of profits. The equilibrium policy function is that estimated in the previous section. We then calculate the simulated counterpart of the value function in equation (1), which involves approximating the distribution of the states considered by the firm. Assuming firms have consistent beliefs about game play—the premise of the BBL method—then a natural approximation of the distribution of states in firms’ consideration sets is the empirical distribution observed in the data, denoted $\hat{S}$. This approximation is also consistent with our policy-function estimation, which matches firms’ behavior to the same distribution of states.

In the final step, we search for structural parameters θ —which here is the vector of managerial costs μ_τ and χ_τ —that rationalize the estimated policy function as the equilibrium one, i.e., the one among the class of admissible policy functions maximizing the value function. Let $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ denote the admissible class of policy functions, comprised by substituting alternative values for the estimated parameters in the monitoring and price-changing equations. Our identification condition is that no deviation in $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ can dominate the estimated policy function $\hat{\sigma}_i$. Formally,

$$E_{s \in \hat{S}}[V_i(s; \hat{\sigma}, \theta)] \geq E_{s \in \hat{S}}[V_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}, \theta)] \quad \text{for all } \tilde{\sigma}_i \in PF(\alpha_\tau, \beta_\tau, C_\tau^k). \quad (14)$$

We estimate θ by finding values that satisfy (14) for a large number of deviations $\tilde{\sigma}_i$. The details of our specific estimator and our choice of deviations are provided in the appendix.

The intuition behind how θ is identified is straightforward. Our estimate of the equilibrium policy function $\hat{\sigma}_i$ will entail some rates of monitoring and price changing. Holding the benefits of these actions constant, their rates should be inversely related to their costs, μ_τ and χ_τ respectively. For example, a high equilibrium monitoring rate entailed by $\hat{\sigma}_i$ must imply a low μ_τ . If instead μ_τ were extremely high, the set of deviations $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ is rich enough that one could be found involving a lower rate of monitoring, reducing the number of times μ_τ is subtracted from the profit stream, which for a given benefit of monitoring would increase the value function. Similarly, a high equilibrium rate of price change (condi-

tional on the rate of monitoring) can only be consistent with a low χ_τ . What may be difficult in practice is distinguishing in the data an equilibrium with frequent monitoring but inert price changing from one with infrequent monitoring but hair-trigger price changing. But that is a difficulty with the identification of the policy function—discussed already in Section 7.2—not the structural parameters. A well-identified policy function will deliver distinct rates of monitoring and price changing that can be used to identify μ_τ and χ_τ . If the policy function is not well identified, then different bootstrapped samples will lead to large swings in the rates of monitoring and price changing, leading to large bootstrapped standard errors on the structural parameters. Thus the standard errors on the structural parameters will serve as a natural check on the identification of the policy function.

To identify the structural parameters, it is crucial that the richness of the set of deviations $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ be exploited. Not all deviations give useful information. For example, a deviation generating negative margins most of the time would be dominated by the equilibrium strategy for any managerial-cost parameters. As discussed further in the next subsection, we will sample from distributions of deviations that have a realistic chance of being profitable to maximize the power of identification assumption (14).

The restriction of deviations to $\tilde{\sigma}_i(s_i) \in PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ is where our procedure accommodates behavioral managers, who pursue rule-of-thumb strategies short of fully rational strategies. A fully rational strategy would have to dominate all conceivable deviations including, for example, a one-time price increase of \$1 in any given hour, a deviation which is not in $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$. Given the difficulty of solving for the fully rational strategies in our setting, we are reluctant to examine deviations that only “work” (i.e., generate inequalities in the correct direction) if firms are fully rational. The restricted set of deviations we consider reflects the idea that firms experiment among a simpler class of pricing rules to discover the most profitable of them. By restricting the set of deviations, our identification assumption is weaker than the standard assumption in BBL, allowing for estimation of structural parameters that is robust to certain forms of behavioral pricing. The only potential pitfall for structural identification would be if the deviation set were not rich enough to span possible combinations of monitoring and price-changing rates. Our model avoids this pitfall, however: fixing all the other parameters, varying the constant terms in α_τ and β_τ over $(-\infty, \infty)^2$ can produce any combination of monitoring and price-changing rates in $(0, 1)^2$. The true drawback to the restriction $\tilde{\sigma}_i \in PF(\alpha_\tau, \beta_\tau, C_\tau^k)$ lies in the possibility of misspecifying equilibrium strategies, hence the importance of establishing good fit of the policy function in Section 7.4.

Additional technical details on the estimation of the structural parameters have been relegated to the appendix, allowing us to turn directly to the results.

8.2. Results

Table 5 presents the estimates of the structural parameters from the profit equation (3). Recall that we rely on estimates of $Base_{it}$ and $Upsell_{it}$ from previous research, leaving managerial costs μ_τ and χ_τ as the only structural parameters to be estimated. We measure μ_τ and χ_τ in dollars.

Type 1 firms are estimated to have a substantial monitoring cost of around \$72, significantly different

from 0 at better than the 5% level. While at first glance this estimate may seem high, further consideration suggests it is plausible. The monitoring cost covers a series of managerial activities including reviewing recent sales and inputs, acquiring competition status, and integrating information from several sources. If these activities occupy about an hour of the manager's time, the monitoring cost should be comparable to the manager's hourly pay. The estimated cost of changing price for a type 1 firm is in fact negative, -6.3 , although insignificantly different from 0. Using the 95% confidence interval, we can reject that costs of changing price exceed \$4.1 at better than the 5% level. The estimates suggest that managers find it costly to continually attend to the market but once they do attend, the menu costs of electronically updating the price are fairly trivial. This estimate confirms the notion that the advent of e-commerce can practically eliminate physical menu costs. The stark contrast between the two types of managerial costs are in align with a series of recent macro papers. For example, Alvarez, Lippi and Paciello (2015) finds the cost of reviewing information is three times as large as the cost of changing price, and Zbaracki, Ritson, Levy, Dutta and Bergen (2004) finds the managerial cost behind the decision of price change is about six times as large as the physical cost of the actual price change.

Turning to type 2 firms, the estimate of the monitoring cost at \$48 is a bit lower though qualitatively similar to that for type 1. The 42.4 estimate for the cost of changing price is substantially higher than that for type 1's. However, the remarkably wide confidence interval $[-177.2, 145.8]$ around the estimate indicates that the cost estimate is essentially uninformative. The small number (eight) of type 2 firms and the infrequency of their price changes combine to generate few price-change observations that go into identifying the policy function. This shows up in standard errors for the policy-function parameters in Table 3 that are around five times higher for type 2 than type 1 firms, leading to the wide confidence interval seen here. The clustering procedure that divided firms up into types allowed us to keep heterogeneous type 2 firms from contaminating the estimates for other types but did not provide enough data to allow credible estimation of that type's managerial costs.

The estimates for type 3 firms are almost identical to those for type 1. They also have a substantial monitoring cost of around \$68 and negative but insignificant cost of changing price of -7.4 . One might have thought the differences between type 1 and type 3 firms stem from differences in managerial capacity. Perhaps type 1 managers occupy the low ranks because they can cheaply monitor and respond to this active segment of the market. In fact, while type 1 and type 3 firms occupy different ranks, they do not differ much in their price-changing behavior, as Figure 4 showed. This similarity naturally translates into similar managerial costs. One possible story is that similar managers are indifferent among various rank positions because high margins and upselling profits (per consumer) at high ranks balance low sales. The indifferent managers are content then to spread and fill out the rank space.

Our preferred estimates given in the bottom of Table 5 re-estimate the structural parameters after imposing a non-negativity constraint on costs. This is a compelling theoretical restriction because it is hard to imagine monitoring or inputting prices providing a utility boost. Perhaps more importantly, the assumption plays an essential role in separately identifying the components of managerial costs in certain cases, in par-

ticular when the policy function happens to generate a price change for almost all monitoring episodes. To understand this essential role, consider the limiting case in which the manager changes price every time he monitors. Then $\mu_\tau = 100$ and $\chi_\tau = 0$ would produce the same net profit as $\mu_\tau = 1,100$ and $\chi_\tau = -1,000$, as indeed would every linear combination of μ_τ and χ_τ on the line including these points. The structural parameters would be unstable and huge positive values for μ_τ and huge negative values for χ_τ could result. Imposing the constraint $\mu_\tau, \chi_\tau \geq 0$ eliminates this instability and selects plausible values of the structural parameters.⁹ The non-negativity constraint binds for type 1 and 3's cost of changing price. That estimate becomes precisely 0, while the cost of monitoring falls slightly for them and its confidence interval tightens. The estimates for type 2 remain unchanged although the 95% confidence interval around that type's monitoring cost now includes 0 and ranges up to 210.9, reinforcing the impression that the structural estimates for that type are simply uninformative.

8.3. Sensitivity

This section gauges the sensitivity of the results to several specification choices made in constructing the profit and value functions. Regarding monetary profits, we adopted a particularly simple specification in equation (3), the markup of price over unit cost times quantity, measuring unit cost as the wholesale acquisition cost of a memory module using our daily data. In practice, firms may have had other unit costs. Credit-card purchases involve a transaction fee, typically 2.25%. There is also a small chance of loss or breakage of the wholesale product. Offsetting these extra costs, firms gained extra revenue through the approximately \$3 difference between actual shipping costs and the \$9.99 they commonly charged the consumers. Our presumption is that these extra costs and revenues largely cancel each other out, leaving equation (3) as an accurate representation of monetary profits. Table 6 shows how the structural parameters change with alternative specifications of monetary profits.

The first row for both the cost of monitoring and price change repeats the baseline estimates from the previous table for comparison. The structural costs presented have all been estimated imposing the non-negativity constraint. Following the baseline estimates, the next row recomputes the structural estimates after adding \$2 to unit cost. In general, higher unit costs lead the estimates of managerial costs to fall. Intuitively, the higher are unit costs, the lower are unit profits; lower managerial costs are required to justify the observed frequency of managerial activity. Qualitatively, though, the results do not change much, as the monitoring cost still dominates the cost of price changing for type 1 and type 3 firms. Indeed, the cost of price changing continues to hit the zero lower bound for type 1 and type 3 firms but it remains positive for type 2 firms.

⁹Although our estimated policy functions happen to generate substantially more monitoring episodes than price changes, for some bootstrapped samples, the estimated policy functions generate similar numbers of monitors and price changes, leading to the problem of wildly positive values of μ_τ and negative values of χ_τ . To address this problem in the unconstrained estimation, for a valid bootstrap draw, we required the number of monitors to exceed the number of price changes by at least 5%. The percentage of valid bootstraps is listed in Tabel 5. A side benefit of imposing the non-negativity constraint is that it eliminates the problem of invalid bootstraps, as indicated by 100% of the bootstrapped sample being valid, as stated at the bottom of the table.

Our specification of monetary profits used the function in Figure 3B for upselling profits. The next row for each managerial cost re-estimates the structural parameter cutting the assumed upselling profits in half for that type. Again, as expected, this leads to lower estimates for managerial costs. For example, monitoring costs for type 1 firms falls from \$68.4 to \$26.0. While this is a large quantitative change, still the qualitative picture remains that for type 1 and type 3 firms, the monitoring cost is substantial while the cost of price changing is negligible. Notably, after this reduction in upselling profits, type 2 firms start to look much like the others. The cost of price changing hits the zero lower bound for them, too. Therefore, we cannot be sure whether the inactivity of type 2 firms is explained by a high cost of changing price or a lack of add-on profits. Some anecdotal evidence for the latter interpretation comes from examining firms' product webpages. While it is quite common for type 1 and type 3 firms to have complex webpages emphasizing available upgrades, most type 2 firms have very simple webpage designs offering little information on premium options.

The last row re-estimates the structural parameters changing the profit specified for firms that disappear from our sample of the first two Pricewatch pages because they rise to rank 25 or higher. As mentioned in footnote 7, this was set to \$4.70, but here we cut it in half. This causes very little change in the estimated managerial cost for type 1 firms because they spend very little time at high ranks. For types-2 and 3, estimated monitoring costs rise to offset the higher benefit from adjusting price to avoid falling off the list. For all three types, the qualitative results are insensitive to this change in specification.

9. Counterfactuals

To assess the importance of managerial costs for the firm's pricing behavior, in this section we present counterfactual exercises in which we shock a single firm i 's managerial costs and see how its pricing behavior and profit change, leaving other firms the same as before. In effect, this section performs the inverse exercise from the structural estimation in the previous section. In the structural estimation, we estimated a policy function from actual pricing behavior and used this to infer firms' managerial costs. Here, we posit a vector of managerial costs for i and search over policy functions for the one maximizing its simulated profits. For each candidate policy function, we want to simulate i 's new pricing behavior and profit assuming that all other firms continue with their originally estimated managerial costs, equivalent to assuming that their pricing behavior is given by the originally estimated policy functions. This is exactly the simulation we did for the 1,800 deviations during the structural estimation, so we can treat these deviations as candidate policy functions, and use the recorded simulations to recalculate simulated profits with the new managerial costs for these deviations, and pick the one with the highest profits as the firm's optimal policy at the new managerial costs. In fact, for this purpose, we did a much finer search by simulating for additional several thousands candidate policy functions that are drawn sequentially from the proximity of the existing promising ones. To save space we will focus on the case in which i is a type 1 firm. The results are shown in Figure 8.

Panel A shows the monthly number of prices changes for the firm as the costs of monitoring and price change vary from 0 to 100 dollars. With a large yet finite pool of high quality candidate policy functions, we

are able to identify 31 policy functions as optimal for some given vectors of managerial costs. The surface is therefore marked by 31 plateaus, which indicate regions of costs for which we identified the same strategy as optimal. Even as costs of monitoring and price changing go to 0, the frequency of price change does not grow without bound. In other words, in the complete absence of frictions associated with price change, it is not optimal for a firm to continuously tweak its price. There are three reasons for this. First, other firms have retained their positive costs, so they do not respond continuously. This results in stretches when state variables do not change during which i 's optimal strategy is to keep price constant. Second, changing price, especially downwards, may trigger other firms' reaction and intensify future competition, so there is a dynamic incentive to wait for a while between price changes. Third, prices in this market are posted in whole dollar amounts, so even if a firm continuously monitored its optimal continuous price, it still would not want to change the price until the optimal price exceeded the threshold necessary to move the price a whole dollar up or down. When i has no costs of monitoring or price changing, it ends up changing price around once a day.

Panel B explores the same counterfactual exercise but now focuses on a different outcome variable: i 's monthly profits. These are the net profits from equation (3), which subtract off the new managerial costs with which we are shocking firm i , accumulated over the month. While the surface in the previous panel had discrete jumps reflecting the discrete changes to firm i 's pricing policy, the continuous changes in i 's costs smooth out its profit function in this panel. We will highlight several key insights that can be drawn from the graph. The graph shows that the division between the two types of managerial costs matters for profit. At the estimated costs of \$68.4 for monitoring and \$0 for price changing, the firm's net profit is \$6,070. Holding the total managerial cost of a joint episode of monitoring and price changing constant at \$68.4 but shifting more of the joint cost from monitoring to price changing, the firm's monthly profits monotonically increase from \$6,070 to \$6,496. In principle, simple arithmetic could explain this gain: holding constant the firm's monitoring and price-changing episodes, it gains from shifting more of the joint cost onto price changing because it monitors multiple times for each price change, so managerial costs would be lower with free monitoring and \$68.4 price changing rather than vice versa. In fact the gain here is due to a more subtle change in the firm's strategy. When monitoring is expensive, the firm ends up forgoing some prime opportunities to change price and changing price is some other less than prime circumstances knowing that it would be too expensive to delay and keep tabs on the market. As monitoring becomes cheaper, the firm can keep almost continual tabs on the market and change price in exactly the right states. At the estimated costs of \$68.4 for monitoring and \$0 for price changing, the firm's optimal strategy leads it to monitor an average of 10.7 times per month, generating managerial costs of $68.4 \cdot 10.7 = 731.9$. If the costs are shifted to \$0 for monitoring and \$68.4 for price changing, the firm's optimal policy now leads it to continually monitor and change price 12.1 times a month, generating managerial costs of $68.4 \cdot 12.1 = 827.6$. Aggregate managerial costs are thus higher with more weight shifted from monitoring to price changing. So it is not an arithmetic reduction in cost that leads to the counterfactual profit increase. Rather this increase in profit comes from improvements in the firm's pricing policy, possible when the managerial activity with additional

option value—monitoring—becomes cheaper even as the managerial activity without option value—price changing—becomes more expensive.

Another key insight from Panel B, which can be drawn from considering the height of the surface, is the potentially large gain to adopting technologies that decrease managerial costs, potentially thousands of dollars a month just for this one product. In particular, based on numbers in this graph, the monthly net profits would increase by over \$1,500 if the managerial costs decreased from the estimated levels to zero. One would underestimate the profits improvement assuming the firm simply saves the managerial cost, which is about a half of \$1,500, without showing a policy response. Re-optimizing the policy function to a suitably active one contributes the other half of the profit improvement. The gain would presumably be multiplied if the technology could be used to reduce managerial costs for the scores of other products the retailers marketed on Pricewatch. In fact the retailers did move to automated pricing soon after the time period of our data, consistent with our estimates of potentially large gains from doing so.

10. Conclusion

In this paper, we studied firms' price-changing behavior in an online market for computer components. Special features of this market made it particularly suitable for study: firms were ranked according to price with lower-price firms receiving more prominent listings and the bulk of the sales on the market; this ranking system, coupled with rapidly changing market conditions, gave firms an incentive to change price frequently as they jockeyed for position. The abundance of price-changing episodes over the year of high-frequency (hourly) observations offers an opportunity to precisely estimate a structural model of price-changing behavior.

We used some initial reduced-form evidence to direct the structural modeling. While firms were free to change price continuously, as frequent as the price changes were, they were far from continuous. Despite competing in a nationally integrated market (at least to some extent), managers on the East and West Coast changed prices at times during the day suiting their shifted schedules, and managers in all locations rarely changed prices on weekends. We took this as reduced-form evidence that the costs of managerial activity provided a source of pricing inertia, motivating a structural model allowing us to separately estimate both the managerial cost of monitoring the market and the managerial cost of entering the new price (a pure menu cost). We were also careful to incorporate firm heterogeneity in the structural model, based on visual examination of price and rank paths for some representative firms showing systematic differences in their strategies, with some firms changing price at least weekly while others only once or twice a month and with some firms aggressively targeting prominent ranks while others higher ranks with lower sales. We incorporated firm heterogeneity in the structural estimates by using a machine-learning method to cluster sample firms method into three types and then estimating policy functions and structural parameters allowed to freely vary across types.

Given the emphasis on frictions in managerial behavior, and the prohibitive complexity of the state

space in our setting, we built a dynamic model of a boundedly rational manager. The manager changes price according to a rule of thumb based on a subset of state variables attended to rather than making the price change that would be optimal given the full state space each instant. The key step for our approach to work is to estimate a rule of thumb—in the language of dynamic structural estimation, a policy function for price changes—that accurately captures managerial behavior. We do this by allowing for a sufficiently rich set of carefully selected state variables, further enriched by allowing for different policy functions across heterogeneous firm types, and then determining that the policy function performs well across a suite of goodness-of-fit exercises. Following BBL, we estimate the structural parameters—here, the managerial costs of monitoring and changing price—as those rationalizing the estimated policy function as being more profitable than deviations. The new feature added to accommodate bounded rationality is that deviations are restricted to the class of admissible rules of thumb rather than any arbitrary price change.

For the types of firms for which we have reliable estimates (types 1 and 3), we estimate a cost of monitoring of roughly \$60 and essentially no cost of price changing. (The estimates for type 2 firms are considerably noisy because there are relatively few of these firms and they changed price infrequently, leaving few price-changing episodes to use to estimate the model.) While these managerial-cost estimates can be moved around by introducing or removing factors in the profit function, the qualitative results remain. Managerial costs primarily arise in the monitoring stage; the further cost of inputting price changes is fairly trivial, suggestive of technological features of this e-commerce market.

This paper fills several gaps in the economics literature. First, we extend methods in the literature on dynamic structural estimation to accommodate behavioral agents who may behave according to rules of thumb based on a subset of state variables. This is an important extension in our setting, where our central focus is on limits to managerial capacity, but may be realistic in other settings as well. Second, we contribute to the empirical macroeconomics literature on retail price stickiness. Like that literature, we suggest that the costs of managerial attention and activity may be important. Our novel contribution is a carefully specified dynamic model of pricing behavior that can generate actual structural estimates of these costs. Finally, we contribute to the behavioral-economics literature by, first, providing a framework for structural estimation in the presence of behavioral agents and, second, by providing numerical estimates of the cost of managerial activity, providing an insight into the psychological barriers of actions such as changing prices that are typically regarded as automatic in neoclassical models.

Appendix A: Details on Structural Estimation

This appendix provides additional technical details on the estimation of the structural parameters omitted from the text.

To transform our identification condition (14) into an estimator of θ requires empirical analogs to the expectations over value functions appearing there, which we will compute via simulation. Our empirical analogue to the first expectation, $E_{s \in \hat{S}}[V_i(s; \hat{\sigma}, \theta)]$, is

$$\widehat{EV}_i(\hat{\sigma}, \theta) = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \sum_{t_{mn}=0}^{\min\{T, 720\}} \delta^{t_{mn}} \pi(s_{t_{mn}}, \hat{\sigma}(s_{t_{mn}}), \theta). \quad (A1)$$

There are three summations in (A1). The first sum simulates the expectation over initial states represented by the $E_{s \in \hat{S}}$ operator. We do this by taking M draws from the set of state vectors observed in the data, $\hat{S}$, and averaging the result (hence the division by M). The second sum simulates the expectation implicit in the value function V_i ; this expectation is over the distribution of all possible histories of game play and private shocks starting from the given initial state. We compute this expectation by simulating N histories for each of the M initial states and averaging the resulting value functions (hence the division by N). The third sum adds up the profit stream implicit in the value function.

The other expectation, $E_{s \in \hat{S}}[V_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}, \theta)]$ from the identification condition (14) is similarly transformed into its empirical analogue $\widehat{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}, \theta)$. Instead of all firms behaving according to the estimated policy function $\hat{\sigma}$, firm i deviates to the policy function $\tilde{\sigma}_i \in PF(\alpha_\tau, \beta_\tau, C_k)$ in the simulation, resulting in profit $\pi(s_{t_{mn}}, \tilde{\sigma}_i(s_{t_{mn}}), \hat{\sigma}_{-i}(s_{t_{mn}}), \theta)$ as the new summand in (A1).

The upper limit in the third sum is modified from the value function as originally appears in (1). Instead of calculating the value function over an indeterminate number of periods ending with firm i 's exit at time T_i , we just add up the profit during the first month—720 hours to be precise. We do this to reduce the accumulation of simulation errors as the period becomes longer. Given the nature of deviations we are considering, with the firm deviating to a whole new policy function for the entire game, profits are stationary in all simulations. Hence the average per-period profit over 720 hours is an unbiased estimate of the average over the whole game. This is not true of some excluded deviations, for example, a one-time increase in price of \$1; such a deviation could generate a complicated impulse response, which would lead the average profit over a truncated period to diverge from that over the full game. We set the discount factor to an annual value of 0.95, though discounting turns out to be fairly inconsequential over the month horizon we are considering.

Calculating $\widehat{EV}_i$ is further expedited following BBL's insight that when the profit function is linear in the structural parameters, this linearity is inherited by $\widehat{EV}_i$ because it is essentially an average over these linear profits. Thus, for example, we can write $\widehat{EV}_i(s; \hat{\sigma}, \theta)$ as the dot product of two vectors

$$\widehat{EV}_i(s; \hat{\sigma}, \theta) = \overrightarrow{EV}_i(s; \hat{\sigma}) \cdot \overrightarrow{\theta}_\tau. \quad (A2)$$

Here $\overrightarrow{EV}_i(s; \hat{\sigma})$ is a vector with four components, corresponding to the four terms in the profit function in (3). The first component is the sum of the stream of firm i 's base profits in a simulation, $\sum_{t_{mn}=0}^{\min\{T, 720\}} \text{Base}_{it_{mn}}$, averaged over the MN simulations, where firms behave according to their estimated policy functions in each simulation. Similarly, the remaining three components are the averages over the MN simulations of, respectively, the total upselling profit over a simulation, $\sum_{t_{mn}=0}^{\min\{T, 720\}} \text{Upsell}_{it_{mn}}$; the number of times firm i monitored market conditions during a simulation, $\sum_{t_{mn}=0}^{\min\{T, 720\}} \text{Monitor}_{it_{mn}}$; and the number of times firm i changed price during a simulation, $\sum_{t_{mn}=0}^{\min\{T, 720\}} 1\{\Delta_{it_{mn}} \neq 0\}$. The second vector in (A2) includes the structural parameters for a firm of type τ : i.e., $\overrightarrow{\theta}_\tau = (1, v_\tau, \mu_\tau, \chi_\tau)$. The linearity in (A2) means that the summary statistics $\overrightarrow{EV}_i(s; \hat{\sigma})$ are all that need to be saved from the MN simulations to later compute $\widehat{EV}_i(s; \hat{\sigma}, \theta)$ for any given

θ because one just needs to take the dot product of the summary statistics and subvectors of the given θ . Hence the simulation and estimation steps can essentially be conducted independently.

The other expectation estimate can be expressed similarly as

$$\widehat{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}, \theta) = \overline{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}) \cdot \vec{\theta}_\tau, \quad (\text{A3})$$

in this case conducting the simulations with firm i using its deviation strategy $\tilde{\sigma}_i$ in the simulation while other firms use their estimated policy functions.

With these estimates of the expectations in identification condition (14) in hand, we can proceed to estimation of the structural parameters θ . Following BBL, define Q to be the change in the value function caused by a deviation in strategy:

$$Q(\hat{\sigma}, \tilde{\sigma}_i, \theta_\tau) = \left[\widehat{EV}_i(s; \hat{\sigma}, \theta) - \widehat{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}, \theta) \right] \quad (\text{A4})$$

$$= \left[\overline{EV}_i(s; \hat{\sigma}) - \overline{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i}) \right] \cdot \vec{\theta}_\tau, \quad (\text{A5})$$

where (A5) follows from (A2) and (A3). The force of identification assumption (14) here is that Q will be non-negative for sufficiently accurate estimates of the expectations and for θ sufficiently close to the true structural parameters. We will estimate θ by assessing a penalty for violations of non-negativity and choosing the value that minimizes the sum of squared penalties over the deviations considered:

$$\hat{\theta} = \underset{\{\vec{\theta}_\tau | \tau=1,2,3\}}{\operatorname{argmin}} \left\{ \sum_{\{\tilde{\sigma}_i | \tau=1,2,3\}} [\min\{Q(\hat{\sigma}, \tilde{\sigma}_i, \theta_\tau), 0\}]^2 \right\}. \quad (\text{A6})$$

For each firm type, we considered a set of 2,000 deviations σ_i from the set of policy functions $PF(\alpha_\tau, \beta_\tau, C_\tau^k)$. Each deviation was generated by adding three perturbations to the type's estimated policy function. The first perturbation involves pure random noise: we multiplied the coefficients $\hat{\alpha}_\tau, \hat{\beta}_\tau$ by log-uniform noise terms and shifted the cut points $\hat{C}_\tau^k$ by uniform noise terms. The second perturbation involves adding correlated noise terms to the constant term in $\hat{\alpha}_\tau$ and the cut points $\hat{C}_\tau^k$, such that the deviations amount to an experiment with changing the frequency of monitoring and frequency of price change conditional on monitoring in opposite directions. The third perturbation is to add correlated noise terms to the cut points $\hat{C}_k$, such that the deviations amount to experiments with trade off between small and large price changes. We chose these perturbations to reflect actual tradeoffs that firms might face in their decision-making. For each deviation $\tilde{\sigma}_i$, we calculate the value function statistics $\overline{EV}_i(s; \tilde{\sigma}_i, \hat{\sigma}_{-i})$ with $M = 10,000$ random draws of initial states with replacement and $N = 1$ simulation for each initial state. For the estimated policy, we calculate the value function statistics $\overline{EV}_i(s; \hat{\sigma})$ with $M = 1,000,000$ random draws of initial states with replacement and $N = 1$ simulation for each initial state to attain greater accuracy.

Because we have chosen a large number of deviations and a large number MN of simulations for each deviation, the second-stage estimate, $\hat{\theta}$, is quite accurate conditional on the first-stage estimate $\hat{\sigma}$. The main source of randomness in $\hat{\theta}$ therefore is the potential error in the first stage. To account for this main source of error, we create 200 bootstrapped samples and repeat the whole estimation procedure. For each bootstrapped sample $n = 1, \dots, 200$, we derive a new estimate, $\hat{\sigma}^{(n)}$, of the policy function. For each n , we carry out the second-stage using $\hat{\sigma}^{(n)}$ in the place of $\hat{\sigma}$ in (A4), (A5) and (A6) and generate new sets of deviations by perturbing $\hat{\sigma}^{(n)}$. The resulting structural estimates $\hat{\theta}^{(n)}$, to minimize the sum of squared penalties represented by the new Q function. Confidence intervals can be constructed from quantiles of the set $\{\hat{\theta}^{(n)} | n = 1, \dots, 200\}$. We could have used bootstrapping to compute standard errors for the policy-function parameters α_τ, β_τ , and C_τ^k reported in Table 3. It turns out that these bootstrapped standard errors are quite close to the standard errors from the maximum-likelihood procedure we chose to report. To simplify the estimation, we did not incorporate uncertainty in $Q(\text{Rank}_{it})$ and $U(\text{Rank}_{it})$ in our standard errors, but we note that those expressions were fairly precisely estimated in Ellison and Ellison (2009a).

References

- Aguirregabiria, Victor and Pedro Mira. (2007) “Sequential Estimation of Dynamic Discrete Games,” *Econometrica* 75: 1–53.
- Alvarez, Fernando E., Francesco Lippi, and Luigi Paciello. (2011) “Optimal Price Setting with Observation Menu Costs,” *Quarterly Journal of Economics* 126: 1909–1960.
- Arbatskaya, Maria and Michael Baye. (2004) “Are Prices ‘Sticky’ Online? Market Structure Effects and Asymmetric Responses to Cost Shocks in Online Mortgage Markets,” *International Journal of Industrial Organization* 22: 1443–1462.
- Artinger, Florian and Gerd Gigerenzer. (2012) “Aspiration-adaption, Price Setting, and the Used Car Market,” mimeo.
- Atkinson, Benjamin. (2009) “Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly, Station-Specific Retail Price Data,” *Energy Journal* 30: 85–110.
- Bajari, Patrick C., Lanier Benkard, and Jonathan Levin. (2007) “Estimating Dynamic Models of Imperfect Competition,” *Econometrica* 75: 1331–1370.
- Barron, John M., Beck A. Taylor, and John R. Umbeck. (2004) “Number of Sellers, Average Prices, and Price Dispersion,” *International Journal of Industrial Organization* 22: 1041–1066.
- Baye, Michael R., John Morgan, and Patrick Scholten. (2004) “Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site,” *Journal of Industrial Economics* 52: 463–496.
- Bils, Mark and Peter Klenow. (2004) “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* 112: 947–985.
- Blinder, Alan, Eile Canetti, David Lebow, and Jeremy Rudd. (1998) *Asking About Prices—A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Bonomo, Marco, Carlos Carvalho, RenÉ Garcia, and Vivian Malta. (2015) “Persistent Monetary Non-neutrality in an Estimated Model with Menu Costs and Partially Costly Information,” Society for Economic Dynamics 2015 meetings paper no. 1339.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert. (1997) “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?” *Quarterly Journal of Economics* 112: 305–339.
- Carlton, Dennis. (1986) “The Rigidity of Prices,” *American Economic Review* 76: 637–658.
- Castanias, Rick and Herb Johnson. (1993) “Gas Wars: Retail Gasoline Price Fluctuations,” *Review of Economics and Statistics* 75: 171–174.
- Cavallo, Alberto and Roberto Rigobon. (2012) “The Distribution of the Size of Price Changes,” National Bureau of Economic Research Working Paper no. w16760.
- Chakrabarti, Rajesh and Barry Scholnick. (2005) “Nominal Rigidities without Literal Menu Costs: Evidence from E-Commerce,” *Economics Letters* 86: 187–191.

- Clay, Karen, Ramayya Krishnan, Eric Wolff, and Danny Fernandes. (2002) "Retail Strategies on the Web: Price and Non-price Competition in the Online Book Industry," *Journal of Industrial Economics* 50: 351–367.
- Cyert, Richard and James March. (1963) *Behavioral Theory of the Firm*. Oxford: Blackwell Publishers.
- Davis, Michael C. and James D. Hamilton. (2004) "Why Are Prices Sticky? The Dynamics of Wholesale Gasoline Prices," *Journal of Money, Credit and Banking* 36: 17–37.
- Doyle, Joseph, Erich Muehlegger, and Krislert Samphantharak. (2010) "Edgeworth Cycles Revisited," *Energy Economics* 32: 651–660.
- Eckert, Andrew and Douglass West. (2004) "Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations," *Journal of Law and Economics* 47: 245–271.
- Edgeworth, Francis. (1925) "The Pure Theory of Monopoly," in *Papers Relating to Political Economy*, vol. 1. London: MacMillan, 111–142.
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. (2011) "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review* 101: 234–262
- Ellison, Glenn and Sara Fisher Ellison. (2009a) "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica* 77: 427–452.
- Ellison, Glenn and Sara Fisher Ellison. (2009b) "Tax Sensitivity and Home State Preferences in Internet Purchasing," *American Economic Journal: Economic Policy* 1: 53–71.
- Goldberg, Pinelopi K. and Rebecca Hellerstein. (2013) "A Structural Approach to Identifying the Sources of Local Currency Price Stability," *Review of Economic Studies* 80: 185–210.
- Golosov, Mikhail and Robert E. Lucas Jr. (2007) "Menu Costs and Phillips Curves," *Journal of Political Economy* 115: 171–199.
- Gorodnichenko, Yuriy and Michael Weber. (2015) "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review* 106: 165–199.
- Harris, Mark N. and Xueyan Zhao. (2007) "A Zero-Inflated Ordered Probit Model, with an Application to Modelling Tobacco Consumption," *Journal of Econometrics* 141: 1073–1099.
- Heckman, James. (1979) "Sample Selection Bias as a Specification Error," *Econometrica* 47: 153–161.
- Hosken, Daniel S., Robert S. McMillan, and Christopher T. Taylor. (2008) "Retail Gasoline Pricing: What Do We Know?" *International Journal of Industrial Organization* 26: 1425–1436.
- Klenow, Peter J. and Olesksiy Kryvtsov. (2008) "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123: 863–904.
- Klenow, Peter J. and Benjamin A. Malin. (2011) "Microeconomic Evidence on Price-Setting," in Benjamin Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, vol. 3. Amsterdam: North-Holland, 231–284.
- Lach, Saul. (2002) "Existence and Persistence of Price Dispersion: An Empirical Analysis," *Review of Economics and Statistics* 84: 433–444.

- Lewis, Matthew. (2008) "Price Dispersion and Competition with Differentiated Sellers," *Journal of Industrial Economics* 56: 654–678.
- Lewis, Matthew. (2009) "Temporary Wholesale Gasoline Price Spikes Have Long-Lasting Retail Effects: The Aftermath of Hurricane Rita," *Journal of Law and Economics* 52: 581–605.
- Lünneman, Patrick and Ladislav Wintr. (2006) "Are Internet Prices Sticky?," ECB Working Paper no. 645.
- Maskin, Eric and Jean Tirole. (1988) "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica* 56: 571–599.
- Means, Gardiner. (1935) "Industrial Prices and Their Relative Inflexibility," U.S. Senate Document 13, 74th Congress, 1st Session.
- Midrigan, Virgiliu. (2011) "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica* 79: 1139–1180.
- Nakamura, Emi and Jón Steinsson. (2008) "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123: 1415–1464.
- Nakamura, Emi and Dawit Zerom. (2010) "Accounting for Incomplete Pass-Through," *Review of Economic Studies* 77: 1192–1230.
- Noel, Michael. (2007a) "Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets," *Review of Economics and Statistics* 89: 324–334.
- Noel, Michael. (2007b) "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market," *Journal of Industrial Economics* 55: 69–92.
- Noel, Michael. (2008) "Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria," *Journal of Economics and Management Strategy* 17: 345–377.
- Pakes, Ariel, Ostrovsky, Michael, and Steven Berry. (2007) "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *Rand Journal of Economics* 38: 373–399.
- Pakes, Ariel, Jack Porter, Katherine Ho, and Joy Ishii. (2015) "Moment Inequalities and Their Application," *Econometrica* 83: 315–334.
- Romesburg, H. Charles. (2004) *Cluster Analysis for Researchers*. North Carolina: Lulu Press.
- Simon, Herbert. (1955) "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics* 69: 99–118.
- Simon, Herbert. (1962) "New Developments in the Theory of the Firm," *American Economic Review* 52: 1–15.
- Slade, Margaret E. (1998) "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices," *Review of Economic Studies* 65: 87–107.
- Stahl, Dale O. (1989) "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review* 79: 700–712.
- Stigler, George J. and James K. Kindahl. (1970) *The Behavior of Administered Prices*. New York: Columbia University Press.

- Terkel, Studs. (1970) *Hard Times: An Oral History of the Great Depression*. New York: Pantheon Books.
- Vuong, Quang H. (1989) "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica* 57: 307–333.
- Wang, Zhongmin. (2009) "(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation," *Journal of Political Economy* 117: 987–1030.
- Zbaracki, Mark J., Mark Tison, Daniel Levy, Shantanu Dutta, and Mark Bergen. (2004) "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics* 86: 514–533.

Table 1: Variable Definitions and Descriptive Statistics

Variable	Definition	Mean	Std. dev.	Min.	Max.	Obs.
Firm-level variables						
<i>Density</i>	Measure of density in price space of firms with nearby ranks	0.60	0.40	0	3	111,276
<i>Margin</i>	Percentage markup over wholesale cost, $100(Price - Cost)/Cost$	1.01	5.65	-20.50	20.38	111,276
<i>NumBump</i>	Net number of ranks bumped since last price change	1.30	3.40	-22	21	111,276
<i>Placement</i>	Placement between adjacent firms in price space	0.58	0.42	0	1	111,276
<i>Price</i>	Current listed price in dollars	70.1	34.6	25.0	131.0	111,276
<i>QuantityBump</i>	Relative change in hourly sales resulting from rank bump	-0.16	0.36	-2.08	2.48	111,276
<i>Rank</i>	Rank of listing in price-sorted order	10.75	6.77	1	24	111,276
<i>RankOne</i>	Indicates whether firm is at rank 1	0.06	0.24	0	1	111,276
<i>SinceChange</i>	Hours since firm last changed price	117.55	146.45	1	1,113	111,276
Market-level variables						
<i>Cost</i>	Wholesale cost	66.40	36.85	23	129	7,740
<i>CostTrend</i>	Trend in <i>Cost</i> over previous two weeks	-0.19	0.71	-2.06	1.53	7,740
<i>CostVol</i>	Volatility of <i>Cost</i> over previous two weeks	1.64	1.08	0.00	4.36	7,740
<i>Night</i>	Indicates hour from midnight to 8 a.m. EST	0.33	0.47	0	1	7,740
<i>Weekend</i>	Indicates Saturday or Sunday	0.29	0.45	0	1	7,740

Notes: Firm-level variables vary over indexes i and t . Market-level variables vary over t . For these variables, descriptive statistics are for the time series of one observation per period.

Table 2: Variable Means by Firm Type

Variable	Combined types (43 firms)	Type 1 (22 firms)	Type 2 (8 firms)	Type 3 (13 firms)
<i>Density</i>	0.60	0.63	0.52	0.56
<i>Margin</i>	1.01	−0.72	1.24	6.20
<i>NumBump</i>	1.30	1.00	2.12	1.64
<i>Placement</i>	0.58	0.58	0.58	0.58
<i>Price</i>	70.06	68.09	85.42	64.90
<i>QuantityBump</i>	−0.16	−0.17	−0.18	−0.11
<i>Rank</i>	10.75	7.78	13.37	18.06
<i>RankOne</i>	0.06	0.09	0.02	0.00
<i>SinceChange</i>	117.55	99.88	255.08	93.33
Observations	111,276	71,460	16,904	22,912

Notes: Shown are means only of variables varying over indexes i and t . Column for combined firms repeats information from Table 1 for comparison.

Table 3: Maximum Likelihood Estimates of Policy Function

Variable	Type 1 firms		Type 2 firms		Type 3 firms	
	Coefficient	Std. err.	Coefficient	Std. err.	Coefficient	Std. err.
Monitoring estimates α_τ						
Constant	-2.25***	(0.10)	-1.71***	(0.54)	-2.83***	(0.25)
Night	-0.68***	(0.05)	-0.84***	(0.23)	-1.30***	(0.21)
Weekend	-0.47***	(0.04)	-0.61***	(0.20)	-0.71***	(0.11)
CostVol	0.04***	(0.01)	0.04	(0.06)	0.05*	(0.03)
CostTrend ⁺	0.13***	(0.04)	0.33**	(0.15)	0.18	(0.12)
CostTrend ⁻	0.09***	(0.03)	0.02	(0.21)	0.24***	(0.07)
QuantityBump ⁺	0.39***	(0.08)	0.69***	(0.26)	0.13	(0.35)
QuantityBump ⁻	0.36***	(0.05)	0.15	(0.26)	0.20	(0.19)
lnSinceChange	0.25***	(0.05)	-0.23	(0.17)	0.47***	(0.12)
(lnSinceChange) ²	-0.05***	(0.01)	0.02	(0.02)	-0.07***	(0.02)
Price change estimates β_τ						
CostTrend	0.11**	(0.05)	0.25	(0.22)	-0.31**	(0.13)
CostChange	0.06***	(0.01)	0.06	(0.04)	0.09***	(0.03)
Margin	-0.02**	(0.01)	0.04	(0.04)	-0.11***	(0.03)
NumBump	-0.05***	(0.02)	0.07	(0.05)	-0.05	(0.04)
Density $\times$ NumBump	-0.10***	(0.04)	-0.28***	(0.08)	-0.06	(0.06)
Placement	0.30***	(0.08)	0.19	(0.32)	0.07	(0.17)
Rank	-0.04***	(0.01)	-0.09**	(0.04)	-0.05**	(0.02)
RankOne	0.62***	(0.11)	α	α	α	α
Cutoff C_τ^{-1}	-0.15	(0.14)	-1.80	(0.48)	-1.62	(0.37)
Cutoff C_τ^1	0.76***	(0.14)	0.55	(0.87)	-0.44	(0.43)
Log likelihood	-6,174.2		-383.0		-1,337.9	
Observations	71,460		16,904		22,912	

Notes: Coefficients from maximum likelihood estimation of equations (6) and (10) separately for each firm type. Heteroskedasticity-robust standard errors clustered by firm reported in parentheses. ^aBecause most of the observations with *RankOne* = 1 are in group 1, equations estimated for groups 2 and 3 constrain *RankOne* coefficient to be the same as estimated for group 1, 0.62. Model includes cutoffs C_k for $k \in \{-5, -4, -3, -2, -1, 1, 2, 3, 4, 5\}$; for space considerations we only report C_{-1} and C_1 . Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 4: Vuong Non-Nested Specification Tests

Model	Versus alternative model		
	B. ZIOP shifting variables	C. ZIOP omitting variables	D. OP omitting monitoring
A. ZIOP preferred specification	3.20	3.92	11.98
B. ZIOP shifting variables		1.82	11.61
C. ZIOP omitting variables			12.03

Notes: Entries are Z statistics from Vuong (1989) specification test comparing non-nested models.

Table 5: Structural Estimates of Cost Parameters

Parameter	Type 1 firms		Type 2 firms		Type 3 firms	
	Estimate	95% c.i.	Estimate	95% c.i.	Estimate	95% c.i.
Without non-negativity constraint						
Cost of monitoring, μ_τ	72.1***	[52.8, 104.2]	48.0***	[11.6, 283.2]	67.7***	[37.4, 133.0]
Cost of changing price, χ_τ	-6.3	[-40.5, 4.1]	42.4	[-177.2, 145.8]	-7.4	[-67.4, 22.2]
Valid bootstraps	100%		92%		93%	
Imposing non-negativity constraint						
Cost of monitoring, μ_τ	68.4***	[50.1, 87.2]	48.0*	[0.0, 207.8]	63.0***	[37.5, 103.8]
Cost of changing price, χ_τ	0.0	[0.0, 4.1]	42.4	[0.0, 145.8]	0.0	[0.0, 22.2]
Valid bootstraps	100%		100%		100%	

Notes: 95% confidence intervals computed via bootstrapping using 200 runs. Bootstrap is taken to be valid if estimated policy function generates at least 5% more monitoring episodes than price changes. Unless a non-negativity constraint is imposed, structural parameters cannot be independently identified and estimates become unstable as the number of price changes converges to the number of monitoring episodes. Statistically significantly different from 0 in a two-tailed test at the * 10% level, ** 5% level, *** 1% level.

Table 6: Sensitivity of the Structural Parameters to Adjustments in the Profit Function

Parameter	Type 1 firms		Type 2 firms		Type 3 firms	
	Estimate	95% c.i.	Estimate	95% c.i.	Estimate	95% c.i.
Cost of monitoring, μ_{τ}						
Baseline estimates	68.4***	[50.1, 87.2]	48.0*	[0.0, 207.8]	63.0***	[37.5, 103.8]
Adding \$2 to unit cost	48.1***	[33.9, 61.6]	41.7	[0.0, 179.7]	56.0***	[32.6, 92.0]
Cutting upselling profits in half	26.0***	[15.3, 38.2]	22.7	[0.0, 130.7]	43.1***	[23.5, 69.3]
Cutting rank 25+ profits in half	71.8***	[53.1, 91.4]	71.7***	[15.4, 282.0]	90.2***	[54.1, 143.4]
Cost of changing price, χ_{τ}						
Baseline estimates	0.0	[0.0, 4.1]	42.4	[0.0, 145.8]	0.0	[0.0, 22.2]
Adding \$2 to unit cost	0.0	[0.0, 2.3]	10.1	[0.0, 114.4]	0.0	[0.0, 218.]
Cutting upselling profits in half	0.0	[0.0, 0.6]	0.0	[0.0, 27.8]	0.0	[0.0, 18.9]
Cutting rank 25+ profits in half	0.0	[0.0, 3.1]	74.0	[0.0, 188.5]	0.0	[0.0, 17.6]


Notes: Baseline estimates are those from Table 5 imposing the non-negativity constraint. All robustness exercises also impose this constraint. Statistically significantly different from 0 in a two-tailed test at the * 10% level, ** 5% level, *** 1% level.

Figure 1: Example Pricewatch Webpage

BRAND	PRODUCT	DESCRIPTION	PRICE	SHIP	DATE/HR	DEALER/PHONE	ST	PART#
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB PC100 SDRAM DIMM - 8ns Gold leads	.- * LIMIT ONE - Easy installation - in stock	\$ 68	9.69 INSURED	10/12/00 12:40:05 AM CST	Computer Craft Inc. 800-487-4910 727-327-7559 Online Ordering	FL	MEM-128-100PCT
Generic	ONLINE ORDERS ONLY - 128MB SDRAM PC100 16x64 168pin	- * LIMIT ONE	\$ 69	INSURED\$9.95	10/11/00 10:59:56 PM CST	Connect Computers 888-277-6287 949-367-0703 Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128MB PC100 SDRAM DIMM	- * LIMIT ONE - InStock, 16x64-Gold Leads	\$ 70	10.75	10/11/00 2:11:00 PM CST	1st Choice Memory 949-888-3810 -- P.O.'s accepted Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128mb True PC100 SDRAM EEPROM DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	- * LIMIT ONE - In stock - with Lifetime Warranty	\$ 72	9.85	10/10/00 11:30:39 AM CST	pcboost.com 800-382-6678 -- P.O.'s accepted Online Ordering	CA	-
Generic	IN STOCK, 128MB PC100 3.3volt unbuffered SDRAM Gold Lead 168 Pin, 7/8ns - with Lifetime warranty	- * LIMIT ONE Not compatible with E Machine	\$ 74	10.95- UPS INSURED	10/11/00 12:44:00 PM CST	Memplus.com 877-918-6767 626-918-6767	CA	- 880060
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB True PC100 SDRAM DIMM - 8ns Gold - warranty	- * LIMIT ONE	\$ 74	10.25	10/9/00 6:53:25 PM CST	Portatech 800-487-1327	CA	-
House Brand	128MB PC100 3.3volt SDRAM 168 Pin, 7/8ns - with LIFETIME WARRANTY	- * LIMIT ONE	\$ 74	10.50 FedEx	10/11/00 10:20:23 AM CST	1st Compu Choice 800-345-8880 800-345-8880	OH	-
Generic	128MB 168Pin TRUE PC100 SDRAM - OEM 16X64	DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	\$ 75	\$10	10/11/00 2:37:00 PM CST	Sunset Marketing, Inc. 800-397-5050 410-626-0211 -- P.O.'s accepted	MD	-
Generic	128MB 16x64 PC100 8ns SDRAM.	- * LIMIT ONE	\$ 77	10.90	10/12/00 9:37:45 AM CST	PC COST 800-877-9442 847-690-0103 Online Ordering	IL	-
Generic	IN STOCK, PC100, 128MB, 168pins DIMM NonECC, - with Lifetime warranty	- * LIMIT 5	\$ 77	\$10.95 & UP For UPS Ground	10/9/00 5:11:10 PM CST	Augustus Technology, Inc 877-468-5181 909-468-1883 Online Ordering	CA	-
Generic	128MB PC100 8NS 16x64 SDRAM - one year warranty	- * LIMIT ONE	\$ 78	Ups Ground \$10.62	10/11/00 5:16:36 PM CST	Computer Super Sale 800-305-4930 847-640-9710 Online Ordering	IL	-
Generic	PRICE FOR ONLINE ORDERS ONLY - PC100 128MB NonBuffered, NonECC 16x64 SDRAM DIMM 3.3V 8ns mem module	- * LIMIT ONE - with lifetime warranty	\$ 78	10.95	10/5/00 6:29:59 PM CST	Jazz Technology USA, LLC 888-485-8872 909-869-8859	CA	ME-GBP100128


Note: Page downloaded October 12, 2002.

Figure 2: Price and Rank Series for Representative Firms of Each Type



Note: Series start later for type 2 and type 3 firms because they entered during the month.


Figure 3: Effect of Rank on Retailer Outcomes




Notes: Panel A is derived from Ellison and Ellison's (2009a) demand estimates, based on their regression of the natural log of quantity on linear rank. Exponentiating yields the equation graphed, $Q(Rank_{it}) = 4.56(1 + Rank_{it})^{-1.29}$. Panel B is derived from Ellison and Ellison's (2009a) estimates of upselling profit. They estimate that a retailer sells an additional $0.97(1 + Rank_{it})^{-0.77}$ units of a medium-quality product with an average markup of \$15.69 and an additional $0.49(1 + Rank_{it})^{-0.51}$ units of a high-quality product with an average markup of \$31.45. The sum of these, for total upselling profit of $U(Rank_{it}) = 15.18(1 + Rank_{it})^{-0.77} + 15.48(1 + Rank_{it})^{-0.51}$, the function graphed.

Figure 4: Distribution of Size and Spell of Price Changes


Panel A: Histograms showing distribution of size of price changes



Panel B: Kernel density estimates of distribution of price spells




Panel C: LOWESS estimates of magnitude of price change as function of spell




Notes: For legibility, histograms in panel A omit bin for zero price change. When this bar is included, the proportions for each firm type sum to 1. Horizontal axis truncated at ± 10 price changes. Densities in panel B estimated using Epanechnikov kernel. Cleveland's (1979) locally weighted regression smoothing (LOWESS) estimated in panel C. A single outlying observation with aspell of over 500 dropped in the graph for type 3 firms. Panel B and C estimated in Stata 14 using default bandwidths.

Figure 5: Price-Changing Activity During the Day for Retailers on Different Coasts




Notes: Graphs show residual probability of price change each hour estimated after partialling out other covariates. So that the probability functions integrate to 1, they have been converted into conditional probabilities, i.e., conditional on a price change occurring during the day. Probabilities computed from estimates from a maximum likelihood model similar to that reported in Table 3 but with *Night* indicator replaced by suite of indicators for Eastern Time hour and interactions between this suite and an indicator for whether the supplier is located on the West Coast. Model estimated on subsample of East and West Coast suppliers only.

Figure 6: Goodness of Fit of Estimated Policy Function by Firm Type




Note: Variables on vertical axis are discounted using same 0.95 annual factor used to compute value functions in the structural estimation. Over the short 720-hour horizon, discounting has a negligible effect on the graphs.

Figure 7: Goodness of Fit of Estimated Policy Function Across Various Initial States



Note: To save space, graphs show results only for type 1 firms. Monetary profit involves just the first two components of π_{it} , $Base_{it} + Upsell_{it}$, not the managerial costs structurally estimated later. Monetary profit for the actual sample is also estimated but estimated based on firms' actual prices as opposed to the simulations, which use prices from the estimated policy function. Naïve forecast assumes firm earns same profit in each of the 720 hours as in the first. Profits discounted using same 0.95 annual factor used to compute value functions in the structural estimation. Over the short 720-hour horizon, discounting has a negligible effect on the graphs.

Figure 8: Counterfactual Scenarios with Various Different Managerial Costs



Note: Counterfactuals show firm's response to a shock to its managerial costs. Dots indicate counterfactuals associated with optimal policies for managerial costs set to the structural estimates. To save space, graphs show results only for type 1 firms. Net profit subtracts managerial costs from π_{it} . Profits discounted using same 0.95 annual factor used to compute value functions in the structural estimation. Over the short 720-hour horizon, discounting has a negligible effect on the graphs.