

An Empirical Study of Pricing Strategies in an Online Market with High-Frequency Price Information

Sara Fisher Ellison
M.I.T. and CESifo

Christopher M. Snyder
Dartmouth College and N.B.E.R.

May 2013

Abstract: We study competition among a score of firms participating in an online market for a commodity computer component. Firms were able to adjust prices continuously; prices determined how the firms were ranked and listed (lowest price listed first), with better ranks contributing to firms' sales. Using a year's worth of hourly data for individual firms, we estimate a model of price adjustment, characterizing the factors driving price changes and measuring how much these factors differ across firms (i.e., strategy heterogeneity). We find evidence of managerial inattention as a cause of the price inertia observed in the market. Using separate estimates of the pricing model for groups of related firms, we simulate counterfactuals involving different mixes of these groups.

JEL Codes: L11, C73, D21, L81

Contact Information: Ellison: Department of Economics, M.I.T., 50 Memorial Drive, Cambridge, MA 02142; tel. (617) 253-3821; fax. (617) 253-1330; email sellison@mit.edu. Snyder: Department of Economics, Dartmouth College, 301 Rockefeller Hall, Hanover, NH 03755; tel. (603) 646-0642, fax. (603) 646-2122, email chris.snyder@dartmouth.edu.

Acknowledgments: The authors are grateful to Hongkai Zhang for superb research assistance and to Patricia Anderson, Michael Baye, Glenn Ellison, Florian Englmaier, Georg Gebhardt, Chakravarthi Narasimhan, Mo Xiao, and various seminar participants for useful comments.

1. Introduction

We study competition among a score of firms selling a common type of memory module directly to consumers through Pricewatch, a price-comparison website. Several features make this market particularly attractive to study. First, it is an example of a market in which firms can make continuous price adjustments, which are immediately transparent to rivals and consumers. Prices determine the order in which the listings are presented to consumers on the website, and this ranking is an important determinant of demand for this relatively homogeneous product. Such markets continue to grow in importance in the Internet era. Second, we have richer data than typically available for academic studies, including information on each supplier's price at hourly intervals and information on suppliers' wholesale acquisition cost, both over the course of an entire year.

After some preliminaries—a literature review in Section 2, background on the market in Section 3, and details on the dataset in Section 4—the heart of the analysis begins in Section 5 with the estimation of a dynamic pricing model. The model allows for behavioral elements to operate through several channels. We specify first-stage equations for the timing of price change, which allows for managerial inattention and other factors that might lead to price inertia. The second-stage equation for the size of price change is not a structural equation derived from maximization of an Euler equation but captures in a reduced-form way whether the average firm is trying to maintain a target margin, a target rank, or both. Although this approach does not deliver structural estimates of parameters, it does deliver estimates of behavioral strategies that can be used to simulate market dynamics.¹ Initially, we keep the estimation procedure simple by running separate regressions for the different stages. Recognizing a possible selection bias with the simple approach, we advance to a full-blown (S, s) model in which firms occupy an interior region of price inactivity unless driven to action in one direction or the other by extreme realizations, estimating all the equations involved using maximum likelihood.

Besides providing evidence on firms' objectives in price setting, whether targeting ranks, margins, or both, we uncover two important features of pricing patterns that are underemphasized in the theoretical

¹In companion research, we are adapting the methods of Aguirregabiria and Mira (2007) and Bajari, Benkard, and Levin (2007) to estimate structural supplier parameters under the assumption of Markov perfect equilibrium. The present approach does not impose the Markov restriction; more broadly, it is agnostic about whether firms are using equilibrium strategies, accommodating the possibility that they are experimenting with different strategies or otherwise learning about the market (e.g., Harrington 1995), or are engaging in low-level strategic thinking (in, e.g., the cognitive hierarchy model of Camerer, Ho, and Chong 2004). The estimates from our pricing model serve a dual role as the policy function in the dynamic structural estimation.

literature on dynamic price competition. The first feature is pricing inertia. Firms do not exercise the option to change prices continuously. Far from it: in fewer than a third of the days in which key market conditions change, either a cost change or a change in rivals' prices large enough to bump the firm in the rankings, does the firm respond by changing its own price. Davis and Hamilton (2004) offer four hypotheses for pricing inertia: menu costs, information-processing lags, customer acceptance, and strategic recognition of rival responses. We find evidence of a fifth: the cost to the manager of continually monitoring market conditions, and computing best responses to these, leading to periods of inattention. We uncover several pieces of evidence for managerial inattention. A firm's price changes tend to be clustered in time, with periods of inactivity even when conditions would merit a price change. Firms are less active when managerial attention is particularly costly: at night and on the weekends. Anchoring time in the online market to Eastern Time, morning starts earlier for firms on the East Coast of the United States than on the West Coast. Concomitantly, pricing activity ramps up earlier on the East Coast compared to the West Coast. Section 6 provides a detailed analysis of this evidence on managerial inattention.

The second feature of pricing patterns we highlight is strategy heterogeneity, discussed in Section 7. By strategy heterogeneity we are not just referring to price dispersion, which is quite wide in this market, but to broader differences in pricing dynamics including whether they maintain target ranks or margins and how far from their targets they are willing to stray before adjusting prices. Such heterogeneity in prices and dynamic strategies is all the more interesting given that the product category we study (generic 128MB PC100 memory modules) exhibits little product differentiation or retailer branding.

We use several approaches to analyze strategy heterogeneity. We conduct a likelihood-ratio test, comparing the likelihood from pooled estimation of our pricing model to that allowing for separate coefficients for each firm, which we have sufficient data over the course of our year sample to estimate, strongly rejecting pooling. In the interest of parsimony, we separate the firms into three strategic groups and re-estimate the pricing model separately for each group. The groups are formed using an agglomeration procedure, iteratively combining groups whose members are closest in terms of the Euclidean distance between vectors of strategic variables, until the final three groups are left. The groups make intuitive sense, with one actively changing price and occupying low-price ranks, another maintaining moderate prices and making few changes, and the third moderately active but charging the highest prices. Our reduced-form approach allows us to document heterogeneity and to flexibly account for it in the estimation, but does not allow us to

identify its origin, whether resulting from cost differences (besides the cost of goods sold, which is captured in our data), differences in managerial competence, or different experiences from experimentation among different strategies or mixing among them.

Section 8 simulates the prices that are chosen by firms using the estimated strategies from the pricing model run separately for the three strategic groups, where the time series of actual costs and a Poisson entry process provide the shocks to the market. The simulations serve two purposes. First, comparing simulated outcomes using the same mix of strategic groups in the data to actual market outcomes provides a specification check on the model. It turns out that the simulations match moments in the actual data reasonably well, so we turn to a second purpose for the simulations, to study counterfactuals. We mainly focus on counterfactuals that allow us to determine the economic importance of strategy heterogeneity, moving beyond mere statistical significance, which can be judged with likelihood-ratio tests. We can see how adding a single firm from different strategic groups impacts the market and see how outcomes would change if all firms started using one of the groups' strategies.

2. Literature Review

We investigate the same market and uses the same data as Ellison and Ellison (2009a, 2009b), but our focus is quite different. Ellison and Ellison (2009a) estimate the elasticity of consumer demand with respect to price, showing that this elasticity can be orders of magnitude higher for items listed on the price-comparison website than for add-on purchases. Ellison and Ellison (2009b) compare price to sales-tax elasticities and document a home-state preference in consumer in their purchasing. The focus in this paper is on the supply rather than the demand side of the market. The demand estimates from Ellison and Ellison (2009a) are used as an input in our analysis.

The two empirical papers closest to ours both in substance and setting are Edelman and Ostrovsky (2010) and Zhang (2010). They document the occurrence of Edgeworth cycles in sponsored-search and online-advertising auctions, respectively.² Those auctions turn out to be similar in design to our price-comparison website. For example, in the sponsored-search auction, advertisers bid on keywords, where a bid is the amount the advertiser pays for every click-through. Unlike a standard auction, bidders besides the highest do not lose the auction; they are simply allocated less desirable spaces, typically lower down on

²In an Edgeworth cycle, studied formally by Maskin and Tirole (1988), firms gradually undercut each other until they reach the zero-profit level where they stay until one relents resulting in a dramatic price rise.

the search-results page. In addition, keyword auctions operate continuously in real time, and advertisers can change their bids at will. Similarly, on the price-comparison website we study, firms continuously “bid” for rank on the webpage through their prices, which can be changed at any time. Gaining favorable rankings requires firms to sacrifice high margins, and indeed winning the rank-1 position often requires the firm to have a negative margin over cost. While we share with Edelman and Ostrovsky (2010) and Zhang (2010) an interest in documenting pricing patterns, our specific findings are quite different from theirs. Prices do not follow Edgeworth cycles in our market. Ranks do cycle, 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. One explanation for the difference is that advertisers in keyword auctions use automated bidding programs, which Edelman and Ostrovsky argue lead to inherently unstable equilibria in a first-price auction format. Prices were set manually in our market, requiring managerial attention. Periods of managerial inattention could lead to pricing inertia and these reverse “Edgeworth cycles” in rank.

Our paper is related to the literature documenting price dispersion in homogeneous, consumer good markets. The markets studied range 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). Perhaps closest to the present paper among these is Baye, Morgan, and Scholten (2004) because they document price dispersion in an price-comparison website, as do we, using rich data collected with a web-scraping technology, similar to ours. However, the focus of our study, firms’ dynamic pricing strategies, is different from theirs.

Our results on pricing inertia and managerial inattention are related to a number of different strands of the literature. Davis and Hamilton (2004) estimate a structural menu-cost model. Their implausibly large estimates of some parameters leads them to reject the model in favor of the most plausible alternative they consider involving strategic rapprochement, although they did not consider an alternative involving managerial inattention for which we find evidence. Other related empirical research includes Lewis (2009), who shows that temporary shocks may have long-lived strategic effects in the wholesale gas market. A theoretical literature providing conditions under which asynchronous pricing emerges in equilibrium includes Fishman (1992), Cahuc and Kempf (1999), and Lau (2001). On limits to managerial processing power, Bhaskar and Vega-Redondo (2002) show that all Nash equilibria must be in Markov strategies if agents have limited memory and strategies conditional on lengthening histories involve complexity costs.

While there is a large literature on price dispersion—so heterogeneity in price levels—there is a smaller literature on heterogeneity in strategies used in a dynamic game. Hosken, McMillan, and Taylor (2008) identify strategy heterogeneity in retail gasoline. Lewis (2008) also documents strategy heterogeneity, finding that his stylized facts are hard to reconcile with a single existing theory. Theoretical papers on strategy heterogeneity include Hansen *et al.* (1996), who analyze a Bertrand duopoly with differentiated products. An unanticipated negative demand shock may elicit asymmetric responses as a decrease in one firm's price reduces the benefit to the other from changing its price, and it may decline to do so if price changes involve a fixed cost. In Lal and Matutes' (1989) model, manufacturers of a product bundle charge different prices for the individual components although the overall bundle is sold for the same total amount. This strategy allows firms to compete for different segments of the heterogeneous population of consumers with their differentiated products. Such strategies may be relevant for our market, because the product under consideration may be offered to induce consumers to buy other products along with it or as a substitute for it, so multiproduct strategies may be important in our market and may contribute to heterogeneity in strategies.

A large literature attempts to characterize interesting dynamics (including Edgeworth cycles) in retail gasoline prices. Besides Hosken, McMillan, and Taylor (2008) and Lewis (2009) cited above, the literature includes Castanias and Johnson (1993), Eckert and West (2004), Noel (2007a, 2007b, 2008), Atkinson (2009), Wang (2009), and Doyle, Muehlegger, and Samphantharak (2010). Our hourly data are collected at a higher frequency than most of these studies, which mostly use daily data, an exception being Atkinson (2009), which uses bi-hourly observations.

Edgeworth cycles are solely a function of firm strategy and not driven by fluctuations in market conditions. Of course, price cycles and other interesting movements may be generated by underlying fluctuations in demand and/or cost. A large empirical macroeconomic literature seeks to understand pricing over the business cycle. See, e.g., Golosov and Lucas (2007), Nakamura and Steinsson (2008), and for a survey, Klenow and Malin (forthcoming). Industrial organization economists have also contributed. For example, Carlton (1986) uncovers a host of stylized facts about price rigidities in manufacturing; Borenstein, Cameron, Gilbert (1997) identify an asymmetry in the response of wholesale gasoline prices to cost rises versus falls. A number of theoretical papers study the impact of market fluctuations in repeated games. Perhaps closest to our market setting are those papers that study the issue in a model that also adds price inertia (in the form of staggered pricing—see Fishman 1992, Eckert 2004, and Leufkens and Peeters 2008).

3. Empirical Setting

We look at the online market for computer components mediated by a price search engine called Pricewatch, the same empirical setting examined in Ellison and Ellison (2009a, 2009b). Readers should consult those papers for a fuller description of the market; here we highlight those features relevant for our analysis of supplier behavior. During the period in which our data were collected, Pricewatch was populated by many small, undifferentiated e-retailers selling memory upgrades, CPUs, and other computer parts. These retailers tended to do little or no advertising, have rudimentary websites, receive no venture capital, and run efficient, profit-maximizing operations. They also tended to receive a large fraction of their customers through Pricewatch. The retailers paid monthly fees to list products, had to abide by Pricewatch rules about transparency of prices, and paid no click-through fees. Potential customers could use Pricewatch to locate a product in one of two ways. They could either type a technical product description, such as “Kingston PC2100 512MB,” into a search box, or they could run through a multilayered menu to select one of a number of predefined product categories, *e.g.*, clicking on “System Memory” and then on “PC100 128MB SDRAM DIMM.” In that case, they would receive back a list of products sorted from cheapest to most expensive in a format with twelve listings per page. These pre-defined categories may contain as many as 350 listings from 100 different websites. Figure 1 contains the first page of a typical list, that for PC100 128MB memory modules from October 12, 2000.

We will be focusing on the dynamic aspects of firm price-setting strategy, and it is worth noting that the Pricewatch lists exhibit substantial turnover from day to day and even from hour to hour. On average, five of the twenty-four retailers on the first two pages of the above-mentioned list will change their prices on a given day. Each price change typically moves several other retailers up or down one place on the list. Some websites are clearly big players that regularly occupy a position near the top of the Pricewatch list. From time to time one may observe a firm sitting in the first position for a week or more, but there is no rigid hierarchy.

Some of this turnover can be attributed to the technology used by Pricewatch and the various websites at the time. Pricewatch is a database-based system which relies on e-retailers’ updating their own prices in its database. Our impression is that almost all of the retailers were setting prices in Pricewatch manually in the time period we study. A typical retailer has dozens or hundreds of products listed in the Pricewatch database, making it impractical to constantly monitor one’s place on each predefined product category and

the current wholesale prices for each product. Instead, a retailer might manually examine its position on the most important Pricewatch categories a few times a day and might look at current wholesale prices at most once or twice a day.

Alternatively, a sudden drop or uptick in sales (which a retailer might notice if he is monitoring incoming sales in real time) could indicate that the retailer has been bumped up or down in rank. Figure 2, derived from demand estimates from Ellison and Ellison (2009a), indicates 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.

4. Data

4.1 Sources

Our data come from Ellison and Ellison (2009a, 2009b). The authors scraped information on the first two pages of listings (12 listings per page for a total of 24) 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 products in this category are, physically at least, relatively homogeneous. 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. The authors took this cost to be common across retailers justified by the fact that the typical retailer made wholesale purchases every afternoon based on that day's retail orders, so little or no inventory was held.

A large number of firms made brief appearances on Pricewatch over our sample period. Since we are interested in the dynamics of firms' pricing patterns, we only retained firms that were present for at least 1,000 hours at some point during the year (approximately one-eighth of our time period) and changed

recorded price at least once (a minimal indication that the listing was not dormant). We also deleted the small number of firms who had multiple products on the first two pages of Pricewatch simultaneously. We were left with 43 firms appearing at some point during the year, with at most 24 present in any particular hour in the sample.

4.2 Variables

Based on these data, we constructed a number of variables that might be factors in firms' decisions about timing and magnitude of price changes. The variables include the firm's rank on the Pricewatch list, its gross margin, quantity, length of time since its previous price change, number of times it had been "bumped" (i.e., had its rank changed involuntarily) since its last price change, and so forth. A description of these variables and summary statistics are provided in Table 1.

Several variables require additional explanation. *Margin* is discussed at length in the next subsection. *Quantity* uses the demand estimates from Ellison and Ellison (2009a) to compute a predicted quantity as a function of the firm's price, the lowest of firms' currently posted prices, and most importantly the firm's rank (this part of the function graphed in Figure 2).³ *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. *Placement* measures the closeness of a firm to its lower-price neighbor as a fraction of the distance between its two closest rivals in price space. For example, if three consecutive firms were charging \$85, \$86, and \$88, the value for *Placement* for the middle firm is $0.33 = (86 - 85) / (88 - 85)$. *Density* measures the crowding of firms in the price space around a particular firm, to be precise, the difference between the prices charged by the next higher-priced firm and the firm three spaces below divided by 4. For example, if five consecutive firms were charging \$84, \$84, \$85, \$86, and \$88, the value for *Density* for the firm charging \$86 is $1 = (88 - 84) / 4$. The definitions of the remaining variables are self-explanatory.

Moving from variable definitions to the descriptive statistics in the table, note that the average price for a PC100 128MB memory module in our sample was about \$67 with a large range, \$21 up to \$131. Most of this variation occurs over time, with prices typically above \$100 at the beginning of the period and down in the \$20 range by the end, mirroring a large decline in the wholesale cost of these modules.

³The precise formula to one decimal is $Quantity = 12.2PriceLow^{1.2}Price^{-3.0}(1 + Rank)^{1.3}$.

4.3 Gross Margins

The variable *Margin* deserves particular attention because it serves as an important factor in our pricing model as well as a gauge of market outcomes. *Margin* equals the difference between the listed price and wholesale acquisition cost of the memory module as a percentage of price, $100(p-c)/p$, referred to as a “gross margin” in accounting. Because they do not reflect other costs of order fulfillment or any fixed costs, gross margins overstate the unit profit earned by the retailer. Despite that, the summary statistics indicate that gross margins in this market are low, on average 4% but some as low as -35%. Several explanations can be offered for these low/negative gross margins. One explanation, the focus of Ellison and Ellison (2009a), is that firms employ an add-on pricing strategy. Another is that firms profit from marking up shipping and handling fees over costs. We will account for both factors in later profit calculations, but it is useful to focus simply on gross margins for now. In particular relative gross margins (either across firms or over time for a given firm) will be informative even if the level of gross margins is an inaccurate measure of unit profit.

In addition to being low, gross margins are surprisingly variable, ranging from -35% to 38%. To help decompose the sources of this variation, Figure 3 plots gross margins over time for two example firms, one tending to occupy the lowest-price ranks (firm A), one the middle ranks (firm B). The figure shows that within-firm gross margins fluctuate widely over time. Firms are clearly not maintaining a constant target for gross margins. Rather, margins fluctuate so wildly due to the combination of volatile costs and pricing inertia. Even though prices change often, these changes are not continuous nor always even daily. Thus a firm does not end up mirroring every cost change with a corresponding change in price. To obtain some measure of this pricing inertia, considering the days in which a firm’s rank or cost changed, it changed price only 31% of the time. A number of possible explanations of pricing inertia can be offered, including stockpiling of inventory, strategic rapprochement, and managerial inattention. We will consider these explanations further in Section 6, providing evidence for managerial inattention.

Figure 3 also illustrates substantial heterogeneity in firms’ pricing strategies. Firm A generally maintains negative gross margins (mean -7.2%, median -7.3%) but firm B positive ones (mean 4.5%, median 4.8%). The two example firms show similar hour-to-hour autocorrelations in margins (0.99 for both), but in the broader sample of firms this autocorrelation varies more widely, ranging down to a low of 0.58. We will study this heterogeneity in strategies further in Section 7.

5. Model of Price Changes

We use a two-stage model of firms' price-setting behavior. The first stage involves the decision made each period (hour) t by each firm i to keep the price from period $t - 1$ or to change it. There will be stretches of time (because of stable market conditions or pricing inertia) during which a firm does not change price. If it does change price in a certain period t , it moves to a second-stage decision about the size (reflecting both magnitude and direction) of the price change. We start with a simple model with separate regressions and a parsimonious set of regressors, working up to an increasingly sophisticated model over a series of subsections.

5.1. Baseline Model

Our baseline model captures the first-stage decision—timing of price change—via two separate probit regressions, one for a price decrease and one for a price increase. This approach has some advantages relative to one that combines price increases and decreases into a single probit for the timing of any price change. Using separate probits, certain variables can have the same directional effect on the probability of price increases and decreases while other variables can have opposite directional effects on price increases and decreases. For example, the probability of a price change in either direction may be low when $Night = 1$. On the other hand, $Margin$ may have differential effects on decreases vs. increases: the firm may most want to decrease price for large positive values of $Margin$ to gain sales; the firm may most want to increase price for large negative values of $Margin$ to avoid losses. All these disparate effects can be captured with a linear specification with separate probits for price increases and decreases. The results are reported in the first two columns of Table 2. The second-stage decision—size of price change—is estimated in the baseline model using a separate ordinary least squares regression in which the dependent variable is the proportional change in price: $\ln(Price_{i,t+1}) - \ln(Price_{i,t})$. The results are reported in the last column of Table 2. Standard errors across all regressions are heteroskedasticity-robust and clustered at the firm level.

Consider the probits for timing of price decreases and increases in columns (1) and (2) of Table 2. The constant terms are large and negative, reflecting the low probability of a price change in either direction in any given hour (0.1% on average). The constant is larger in (1) than (2), reflecting the general trend of falling prices in the market.

The coefficient on $Margin$ is not statistically different from 0 in column (1). While one would certainly

not expect firms to respond to large negative margins by decreasing price further, one might expect firms to try to increase sales in the face of large positive margins by decreasing price. The result, which is a fairly precisely estimated 0, indicates that firms are generally content to earn positive margins while they can and not change price then. By contrast, the coefficient on *Margin* in (2) is significantly negative at better than the 1% level, consistent with firms' strong desire to avoid losses by increasing price in the face of large negative margins.

The coefficient on *NumBump* is highly statistically significant in both columns, positive in (1) and negative in (2). These results are consistent with firms having target ranks, reducing price if they have been bumped up from their target and increasing price if their rank has fallen below their target. The coefficient on *CostTrend* indicates that firms are more likely to reduce prices in the presence of a falling cost trend and to increase prices in the presence of a rising cost trend, although the latter effect is not statistically significant. Unpredictable movements in costs captured by *CostVol* has no measurable effect on price decreases but has a significant effect of increasing the probability of a price increase, consistent with firms' desire to pad their margins as insurance against losses if costs plunge unexpectedly.

The coefficient on *SinceChange* is of identical size and significance in columns (1) and (2). The negative value indicates that the hazard of changing price is decreasing over time. Thus, rather than changing prices on a systematic schedule, firms' price changes tend to be bunched in time, followed by longer spells of inactivity than would be predicted by regular schedule. This finding is consistent with a managerial inattention model, but of a certain type, with the manager drawn away from watching the market at random times for random durations, but then making rapid adjustments in periods he or she is allowed to attend to the market, as opposed to setting prices and then returning to revise them at consistent intervals, say once every other day.

The last two variables, *Night* and *Weekend*, contribute more to the timing of price changes than any of the other covariates. The coefficients are negative, of large magnitude, and highly statistically significant. Expressed as marginal effects, the probability of a price decrease is over five times greater and of a price increase is over nine times greater during the day than at night. The probability of a price decrease is over three times greater and of a price increase is over 2.5 times greater during the week than over the weekend. These findings are supportive of an inattention story, where the cost of a manager's attentiveness increases

during nights and weekends.⁴

The last column of Table 2 presents the results from the linear regression on the size of price change. These regressions were run on only the subset of observations for which there was a price change from the previous hour. This results in a much smaller sample because price changed in fewer than 1% of the observations. The dependent variable is the change in log price, $\Delta_{it} = \ln(\text{Price}_{i,t+1}) - \ln(\text{Price}_{i,t})$, essentially a proportional change for small price changes. As indicated by the time subscripts, this is a forward price change from time t to $t + 1$; we regress this on values of the covariates as of time t . Looking at overall fit, the R^2 in column (3) is several times higher than the pseudo R^2 in columns (1) and (2), manifesting the ease of predicting the size of a given price change relative to predicting its exact timing at an hourly frequency.

The coefficients all have the expected signs. The constant term is negative and significant, indicating that prices were typically adjusted downward over the sample period. *Margin* is also negative and significant. This result is consistent with firms' attempting to maintain a gross-margin target, i.e., they increase price when margins become too negative and decrease price when margins become too positive.⁵ *NumBump* is negative and highly significant, consistent with firms' attempting to maintain a target rank.⁶ The magnitude of the coefficient implies that being bumped up one position in the ranking leads a firm to decrease price more than half a percent (0.63%). (Symmetrically, being bumped down one position leads to a 0.63% increase in price.)

CostTrend is not a statistically significant determinant of the size of price change. It may be that the trend reflects the predictable part of past cost, which was already incorporated in the size of the previous price change. Instead, the unpredictable part of cost changes captured by *CostVol* is a significant determinant. The positive coefficient is consistent with firms' maintaining a bigger cushion for margins when costs are less predictable, perhaps to avoid large negative margins without having to change price frequently. In the next subsection, we will see that moving to more refined specifications preserves the signs but reverses the significance levels of the cost variables, so we will avoid placing undue weight on these interpretations.

⁴If demand fell to zero during nights and weekends, it would be impossible to separate an inattention story from a story where a perfectly attentive manager did not change prices because his actions would have no effect. Demand, however, does not fall to zero during these periods. See Figure 2 in Ellison and Ellison (2009b).

⁵To be more precise, the model suggests that the target gross margin is slightly negative. This can be seen by plotting the fitted values of the dependent variable as a function of *Margin* evaluated at the sample means of the other covariates. The horizontal intercept of the resulting line, which is the model's prediction for the gross-margin target, is $\text{Margin} = -7.7\%$. The discussion of negative gross margins in Section 4.3 provides a rationalization for a negative target.

⁶Similar to the exercise in the previous footnote, we can plot the fitted values of the dependent variable now as a function of *NumBump* evaluated at the sample means of the other covariates. The horizontal intercept of the resulting line is essentially $\text{NumBump} = 0$, suggesting that the firm prefers to maintain the status-quo price if it has not been bumped.

5.2. Joint Maximum Likelihood Estimation

Using three separate regressions to estimate a pricing model has the advantage of being straightforward but has two associated disadvantages. First, because the probits for price decreases and increases are independent, when we later use these equations to simulate a price path, it is mathematically possible that a firm is predicted to simultaneously decrease and increase price, a contradiction. This is not a major disadvantage because such instances would be rare. A more significant disadvantage is possible selection bias in the size of price increase equation in column (3) of Table 2. Selection bias may arise if firms only register the true size of their desired price changes if this is either a large negative or large positive change. Under typical/moderate conditions, the firm may be inattentive or may not bother to register its desired price change.

The bias from running a least squares regression either using only the observations for which there was a non-zero price change or including all the observations with zero used for the dependent variable for observations without a price change is similar to that familiar from Heckman (1976). The difference from Heckman in our setting is that rather than having a single threshold separating an inactive from an active region, in our setting selection may occur at two thresholds, a lower threshold below which the firm wants to decrease price and an upper threshold above which the firm wants to increase price, with inactivity observed in an interior interval, i.e., prices may be adjusted according to an (S, s) model.⁷

Figure 4 provides an example of how selection can operate in our setting. The covariate graphed on the horizontal axis has a positive relationship with desired price change, graphed on the vertical axis. The dotted line shows the true relationship around which observations given by circles are scattered. The open circles are desired price changes that fall between the thresholds and are thus observed as zero actual price changes. The solid circles fall outside the thresholds and are thus the only non-zero price changes actually observed. The solid line is the least-squares regression run with just the solid circles. Notice that its slope differs from the true line, illustrating the selection bias. In particular, selection on high negative and positive errors leads to a pivoting of the regression line to become steeper than the true line. In other words, selection in an (S, s) model tends least-squares coefficients to be biased away from 0. This is the opposite of the bias typically found in settings with a single selection threshold. For example, in that setting, Greene (1981) shows that

⁷See Barro (1972) for the first theoretical application of an (S, s) models to price setting with adjustment costs and Golosov and Lucas (2007) for a discussion of selection effects in (S, s) models of price setting.

least squares coefficients are biased toward zero when covariates and errors are normally distributed.

We will address the problems of simultaneous price decreases and increases as well as selection by estimating a model of the timing and size of price change jointly along the lines of an (S,s) model in Figure 4. Let

$$L_{it} = W_{it}\theta + e_{it} \quad (1)$$

be firm i 's latent desire to increase price between time t and $t+1$ and

$$\Delta_{it}^* = X_{it}\beta + u_{it} \quad (2)$$

be the size of the desired price increase, where W_{it} and X_{it} are row vectors of covariates, θ and β are column vectors of parameters, and e_{it} and u_{it} are mean-zero bivariate normal errors with variances 1 and σ , respectively, and correlation coefficient ρ . The actual (as distinct from desired) size of price increase Δ_{it} depends on the value of L_{it} relative to lower and upper thresholds. If L_{it} falls below a lower threshold T_{it}^l , then the firm decreases price: $\Delta_{it} = \Delta_{it}^* < 0$. If L_{it} rises above an upper threshold T_{it}^u , then the firm increases price: $\Delta_{it} = \Delta_{it}^* > 0$. If L_{it} is between the thresholds, despite a possibly non-zero desired price change Δ_{it}^* assuming it were costless for it to attend to market conditions and adjust price, the firm does not change its actual price: $\Delta_{it} = 0$. We will model the thresholds as a function of covariates related to a firm's attention cost as well as factors that might have a non-monotonic effect on L_{it} in which case it would not suffice to have them appear linearly in W_{it} . Formally, let

$$T_{it}^l = Z_{it}\gamma^l \quad (3)$$

$$T_{it}^u = Z_{it}\gamma^u, \quad (4)$$

where Z_{it} is a row vector of covariates and γ^l and γ^u are column vectors of parameters. The same covariates appear in both thresholds because the factors that influence price decreases could also influence price increases in principle. However the parameters γ^l and γ^u can be different in magnitude and sign. Figure 5 provides a schematic diagram for the joint model. Statistical error is assumed to already be incorporated in e_{it} , so no separate errors appear in the thresholds.

Under the assumption of bivariate normal errors, the model can be estimated jointly using maximum

likelihood. The covariates for W_{it} , Z_{it} , and X_{it} were chosen to make the model as comparable to that in Table 2 as possible. In particular, we use the covariates in columns (1) and (2) that can take on positive and negative values (*Margin*, *NumBump*, and *CostTrend*) in W_{it} because the lack of lower bound on these covariates at 0 will allow the index $W_{it}\theta$ in principle to drive both price decreases and increases. The remaining variables that only take on non-negative are included in Z_{it} . Because there are no joint restrictions on coefficient vectors γ^l and γ^u , the Z_{it} can make price increases and decreases both more probable, both less probable, or have opposite effects on the probability of each. The covariates from the least-squares regression of the size of price increase in column (3) of Table 2 are used for X_{it} .

Table 3 reports the parameter estimates $\hat{\theta}$, $\hat{\gamma}^l$, $\hat{\gamma}^u$, and $\hat{\beta}$ in that order. Given the choice of covariates, the results are directly comparable to those in the previous table. It can be shown that the results from column (1), with the signs adjusted by multiplying by -1 , combined with the results from column (2) provide estimates of the same parameters as column (1) of the previous table.⁸ Both are consistent because there is no issue of selection bias with these equations. Similarly, it can be shown that the results from column (3) of Table 3, with the signs adjusted by -1 , combined with the results from column (1) in that table provide estimates of the same parameters as column (2) of the previous table. By contrast, while the covariates in column (4) of Table 3 are the same as in column (3) of Table 2, but the estimates need not be close to each other because of the selection bias inherent in the latter.

Turning to the results from the joint maximum likelihood model in Table 3, as expected, the coefficients in columns (1), (2), and (3) are close to their (sign-adjusted) analogues in columns (1) and (2) of Table 2, in many cases differing only in the third significant digit. The covariates in column (1) have strong and statistically significant effect on the latent desire to increase price in the expected direction. Large negative margins increase the latent desire as does being bumped down in the ranks toward the first position, as does a rising cost trend. The covariates appearing in the thresholds in columns (2) and (3) also have strong and significant effects in intuitive directions. The absolute value of the constant is higher in (3) than in (2),

⁸To see this result, note that a price decrease is observed in the joint maximum likelihood model if $L_{it} < T_{it}^l$. Substituting from (1) and (3) and rearranging, the probability of this event is

$$\Pr(e_{it} < Z_{it}\gamma^l - W_{it}\theta) = \Phi(Z_{it}\gamma^l - W_{it}\theta) \quad (5)$$

$$= \Phi\left([Z_{it} \parallel W_{it}][\gamma^l \parallel (-\theta)]\right), \quad (6)$$

where Φ is the standard normal distribution function and \parallel is the vector concatenation operator. By equation (6), consistent estimates of γ^l and $-\theta$ can be obtained by running a probit of an indicator for a price decrease on the stacked vector $Z_{it} \parallel W_{it}$ as done in column (1) of Table 2. In Table 3, column (1) provides a consistent estimate of θ and column (2) of γ^l .

indicating that price increases are rarer than decreases. The remaining variables operate on the thresholds to expand or contract the region of inactivity for price changes. *SinceChange*, *Night*, and *Weekend* move the lower threshold down and the upper threshold up, expanding the range of inactivity, reducing the probability of both price increases and decreases. Higher *CostVol* moves the thresholds in toward 0, increasing the probability of both sorts of price changes. Firms seem to attend less to prices during the night and weekend, attend less during a spell of inattention, but attend more when conditions are volatile.

The results for desired size of price increase in column (4) generally have the same sign as the analogous results in column (3) of the previous table but have different magnitudes. Consistent with the illustration in Figure 4 that selection in an (S,s) model will tend to bias least-squares coefficients away from 0, we see that the coefficients fall in magnitude in the joint estimation which corrects for selection. For three of the four covariates, the magnitude of the coefficients are cut about in half. A formal test of selection bias can be obtained by comparing the likelihood from the unrestricted joint model reported in Table 3 to that from a model restricting correlation coefficient ρ to be 0. If $\rho = 0$, the size of price increase conditional on a change is restricted to be independent of the desire to change. The resulting estimates in column (4) of Table 3 after imposing the restriction are identical to the least-squares results from column (3) of Table 2. The likelihood-ratio test statistic, which has a χ^2 distribution with one degree of freedom under the null, is a huge 868.2, allowing rejection of the null of no selection at better than the 1% level. This provides strong justification for moving from separate regressions to the joint maximum likelihood model.

The estimates from column (4) of Table 3 generally conform to expectations. The one exception is the negative and significant coefficient on *CostTrend*, implying that a rising cost trend leads to smaller desired price increases. The result is less surprising when combined with the earlier result in column (1) that a rising *CostTrend* leads to more frequent price increases. Firms appear to make more refined adjustments in the presence of strong cost trends.

5.3. Expanding Set of Covariates

The model in Table 3 captures a firm's response to rivals' behavior in a parsimonious way: how far rivals' price changes have moved the firm from its target rank, measured by *NumBump*. We adopt this as our baseline model because richer controls exacerbate problems of over-determination; i.e., it becomes difficult to identify whether a firm is responding to a change in market conditions or a rival's response to the change

in market conditions. Still, it is interesting to bring in more of these strategic considerations. To this end, Table 4 adds a rich set of covariates capturing rivals' entry and the distribution of rivals' prices (indicated by (a)) to the baseline model.

The first added variable, *Quantity*, is mainly a function of rivals' prices through the indirect channel of its dependence on rank.⁹ The positive coefficient in column (1) implies that the firm is more likely to decrease price when quantity is low and increase price when quantity is high, although conditional on changing price, quantity does not influence the desired size of price increase, natural because quantity is in levels and the price change is in percentage terms. The coefficients on *Placement* can be interpreted as saying that the firm is less likely to decrease price and chooses a smaller increase the closer it is to its lower-priced neighbor. One might think that opposite result should hold as the firm would profit from the improvement in rank gained from undercutting this lower-priced neighbor. However, the closeness in price space may induce strategic rapprochement between them. Jumping down the list of covariates to *Density*, a denser local price space is associated with a lower threshold for price increases, thus more frequent price increases, but conditional on an increase the size is smaller. Firms appear to make more frequent tweaks to price in a dense price space.

The next three variables capture the shape of the distribution of market prices and margins in a parsimonious way, here reflecting the minimum and range of prices and the mean of margins in the market. They are difficult to interpret individually because changing one variable while keeping the other two constant requires a complicated reconfiguration of market prices and costs. Overall, each is statistically significant in at least one equation as is the combination of them across equations. The likelihood-ratio test statistic, which has a χ^2 distribution with three degrees of freedom under the null, is 83.9, allowing rejection of the null of exclusion at better than the 1% level. However, their addition does not drastically change any of the coefficients on the other covariates.

The last variable, *NewFirm* is negative and significant in column (1), indicating that a firm's price decreases are more frequent and price increases less frequent after an entrant replaces an incumbent on the first two Pricewatch pages. Offsetting the greater frequency of price decreases, the positive coefficient in column (4) indicates that the magnitude of the price decreases is smaller after an entry episode. Together, the results suggest that entry is associated with existing firms' paying greater attention to fine tune their prices without appreciably affecting their price levels. The explanation need not be causal: periods in which entry

⁹See the formula for *Quantity* in footnote 3. As the formula shows, *Quantity* also depends on rivals' prices through the inclusion of *PriceLow*.

is attractive may be ones in which incumbents have a greater payoff from attending to pricing decisions.¹⁰

6. Managerial Inattention

The results so far have provided some evidence that pricing inertia is due to managerial inattention. The coefficients on *SinceChange* in Table 3 show that the region of inactivity expands as time passes since the last price change, implying that firms tend to have spells of inactivity in pricing rather than changing them at systematic intervals. The coefficients on *Night* and *Weekend* also show an expansion of the region of inactivity during those periods. When it is costly for a manager to attend to pricing, the manager is less attentive.

A finer test can be provided by recognizing that while firms sell to an integrated, national, online Price-watch market,¹¹ they are headquartered in different time zones. Relative to some fixed market time (we use Eastern Time for these purposes), night falls three hours earlier for firms on the East than on the West Coast. Hence the cost of attending should vary over market time differently for East and West Coast suppliers. To test this, we restricted the sample to firms headquartered on the East and West Coast (more precisely Eastern and Western Time Zones). We replaced *Night* in the threshold equations in Table 3 with fixed hour effects and the interaction of the fixed hour effects with an indicator for West-Coast firms.

The results focusing on the newly included variables are reported in Figure 6. The upper panel shows the probability of a price change for suppliers on each coast as a function of the hour of day h estimated in the joint model. The unconditional probability functions are difficult to compare because they have different integrals. To make them more comparable, we converted them using Bayes Rule into conditional probabilities—i.e., the probability of a price change in hour h conditional on exactly one price change during the day. These conditional probabilities have proper density functions that integrate to 1. The notes to the

¹⁰A causal interpretation would require entrants to behave differently than the incumbents they replace. To provide evidence for this hypothesis, we ran a fixed-effects regression of *Rank* on *Tenure* for the subsample of firms entering after the start of our sample (the starting date is left-censored for firms in the market at the start of the sample), including fixed firm and day effects. (To reduce the number of time fixed effects to estimate, we collapsed the dataset down to one hourly observation per day and included day rather than hour effects.) The coefficient on tenure was -0.016 with a standard error of 0.005 , statistically significantly different from 0 at better than the 1% level. The result implies that a firm's rank tends to drift down 1.6 positions for every 100 days it remains in the market, thus implying that new firms tend to enter at the higher ranks. It is not obvious why entry at high ranks would cause incumbents to be more attentive to pricing. We also tested to see whether a firm was more or less active with price changes with tenure by adding tenure as a covariate in the threshold equations in column (2) and (3) of Table 4. Again, we restricted the sample to only those firms entering after the start of our sample. The result was insignificantly different from 0. Thus, there is no evidence that entrants' greater attention to pricing causes incumbents to attend more.

¹¹More precisely, Ellison and Ellison (2009b) find a statistically significant but small effect of shipping times: waiting the four days for the good to be shipped across the country estimated to have about the disutility of a 50-cent price difference. They find an additional home-state preference equivalent to about a \$2 price difference.

figure provide more detail behind the calculations. Just as local time shifts as one moves from East to West Coast, pricing activity shifts as well. Pricing activity by West Coast suppliers does not start until later in Eastern Time compared to East Coast suppliers. Interestingly, the graph for West Coast firms is not merely a three-hour shift of the graph for East Coast firms. The peak comes at 10 a.m. Eastern Time for East Coast firms but not until ten hours later, 8 p.m. Eastern Time, for West Coast firms. It appears as if East Coast managers generally start their day by attending to prices, while West Coast managers end theirs that way.

The concept of first order stochastic dominance captures formally the intuitive idea that pricing activity happens later for West Coast than East Coast firms. The literature is fairly divided on appropriate statistical tests of first order stochastic dominance even with respect to which hypothesis should be the null and which the alternative (see Davidson and Duclos 2000 for a survey). We offer a new combination of tests that best tests the concept in our setting. The tests work by comparing the cumulative distribution functions (cdfs) shown in panel B. We will conclude that the West Coast cdf first order stochastically dominates the East Coast cdf if (a) the area between the East and West Coast cdfs is significantly positive and (b) there is no point h for which the West Coast cdf is significantly greater than the East Coast cdf.¹² The test of (b) is trivial in our case: the West Coast cdf is never above the East Coast cdf, so cannot be statistically significantly greater. To test (a), we derive the analytical expression for the integral between the cdfs, which is a nonlinear function of the parameters that can be tested for difference from 0 using the delta method.¹³ The resulting test statistic, which has a χ^2 distribution with one degree of freedom under the null, is 30.4, statistically significantly different from 0 at better than the 1% level.

7. Strategy Heterogeneity

Although to this point we have presented analysis pooling all firms to get a sense of behavior on average, we have already seen suggestive evidence of strategy heterogeneity as early as in Figure 3. In that example, one

¹²A number of papers suggest tests similar to (b) (see e.g., Bishop, Formby, and Thistle 1992). We have the advantage in our setting that the grid of points at which the tests are undertaken is not arbitrary but determined by the discrete nature of our hourly data. Further, this test alone does not provide the desired evidence that the West Coast cdf stochastically dominates the East Coast, only that the West Coast is not stochastically dominated by the East Coast. Our contribution is to combine (b) with (a). The test we suggest in (a) is to our knowledge novel in the literature. Among other alternatives, McFadden (1989) tests for a difference at the point of maximum difference between the cdfs. Our approach of considering the integral between the cdfs is a more direct measure of stochastic dominance and can be analytically computed in our setting because the random variable (h) is discrete.

¹³The cdf for location $l \in \{E, W\}$ is $F_l(h) = \sum_{j=1}^h \delta_{lj}$, where δ_{lj} is the probability derived in the notes to Figure 6. Because h is discrete in our setting, the integral between the cdfs is actually a sum: $\sum_{h=6 \text{ a.m.}}^{2 \text{ a.m.}} [F_E(h) - F_W(h)]$. Substituting the previous expression for the cdfs into this sum and rearranging yields the formula for the test statistic: $\sum_{h=6 \text{ a.m.}}^{2 \text{ a.m.}} (24 - h + 1)(\delta_{Eh} - \delta_{Wh})$.

of the firms had generally low (indeed negative) gross margins and the other much higher. We noted in the discussion of the figure that while these two firms exhibited similar hour-to-hour autocorrelation in margins (0.99), this autocorrelation varies greatly in the wider sample of firms, down to a low of 0.58.

To test strategy heterogeneity more formally, we returned to the baseline model in Table 3 but allowed all the parameters to differ by firm. For 60% of the firms (26 of 43), there are enough time-series observations to estimate firm-specific parameters. Restricting the sample to these 26 firms, we compared the likelihood from the pooled model to that with firm-specific coefficients. The resulting likelihood-ratio test statistic, distributed χ^2 with 500 degrees of freedom, is 1,615.7, significantly different from 0 at better than the 1% level.¹⁴

To accommodate strategy heterogeneity in our subsequent empirical analysis, we will adopt a compromise between full pooling and 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 categorized firms into a small number of strategic classes and allowed the model parameters to differ freely across the classes.

A number of methods can be used to partition the firms into classes. The method that we found produced the best results for the simulations discussed in Section 8, is referred to as “cluster analysis”.¹⁵ A standard reference for cluster analysis, which is better known in social sciences outside of economics, is Romesburg (2004). To avoid any confusion, 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).

For the particular variant of cluster analysis we used, 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

¹⁴To calculate the degrees of freedom, the baseline model has 20 parameters—the 18 coefficients as well as σ , and ρ —leading to $20 \times 26 = 520$ parameters in the unrestricted model and 20 in the pooled model for a difference of 500. The threshold for 1% significance for a χ^2 random variable with 500 degrees of freedom is 429.4.

¹⁵We also used an alternative method which starts from the strategic classification from the cluster analysis, but then iteratively reallocated firms to groups when this resulted in an improvement in the sum of the log likelihoods, iterating until no further likelihood improvements could be made. Estimating the pricing model using this grouping and then simulating gross margins and profits, the simulation averages are farther from actual averages in the data than the simulations based on the original strategic classification.

the firms' control rather than outcomes depending on external factors.¹⁶ The next step is to standardize the variables 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.¹⁷ We iterated until three strategic groups were left, which we thought achieved a balance between spanning most of the important firm heterogeneity and still obtaining strategic groups large enough to analyze empirically with some confidence.

Table 5 describes the three strategic groups resulting from the procedure. The 22 firms in strategic group 1 generally occupy the lower price ranks and change prices relatively frequently. The eight firms in strategic group 2 occupy the middle price ranks and tend to change prices infrequently. The 13 firms in strategic group 3 generally charge the highest prices but are more active in changing prices than cluster 2. This characterization is echoed in Figure 7, which graphs prices and ranks for a representative firm from each strategic group for a representative month. As is typical in our sample, the representative month was marked by falling wholesale prices, so a firm that maintained a constant price would tend to get bumped higher in rank as firms around it were lowering their prices in response to falling costs. As the means from Table 5 suggest, the price of the group-2 firm is between the low price of the group-1 firm and the high price of the group-3 firm. It is noticeably less active than either of its counterparts, changing price only twice during the month. As a result of its inaction, it is bumped up from a rank of 6 at the beginning of the month to 14 by the end. Both the group-1 and group-3 firms are close to their initial rank at the end of the month.

Table 6 presents the estimates from the baseline model analogous to Table 3 but allowing the parameters to vary across the three strategic groups. The coefficients are qualitatively similar in signs and significance levels but differ in magnitudes. The statistic from the likelihood ratio test of equality of parameters (coefficients and variance parameters) across the three groups is 276.1, distributed χ^2 with 40 degrees of freedom, significantly different from 0 at better than the 1% level. One of the more interesting differences regards the groups' responses to cost volatility. All groups respond to an increase in *CostVol* by shifting their thresholds

¹⁶The 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*.

¹⁷The 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. We used Stata's cluster analysis command to perform the clustering, and these are standard options.

for price changes in toward 0. With groups 1 and 3, *CostVol* does not appear to have a further effect on the size of their price increases. The coefficient of 0.0111 in the last column for Group 2 suggests that they build in a greater margin when costs are uncertain as insurance against cost increases.

8. Simulations

In this section, we report on market simulations where we have firms set prices according to the model estimated in Table 6. The simulations serve two purposes. First, comparing simulated outcomes using the same mix of strategic groups in the data to actual market outcomes provides a specification check on the model. The baseline model, because it includes a parsimonious set of variables and is not built on a first-order condition from a profit-maximization problem, may be missing key determinants of price changes. On the other hand, the baseline model may “overfit” the data by including costs in various forms directly as well as indirectly through rivals’ responses to costs, which can move the firm around in the ranks, affecting *NumBump*. The clustering procedure produced strategic groups with some internal logic, but the procedure was not based on optimizing fit and three groups may not be adequate to capture all the strategy heterogeneity in the market.

As it will turn out that the simulations match moments in the actual data reasonably well, a second purpose for which we will use them is to simulate counterfactuals. We will consider several sorts of counterfactuals including adding one firm of a given type and changing the mix of strategic groups so that all of them are of a given type. The counterfactual simulations play a special role in our analysis which merits some discussion. Our intent is to better understand the nature of the behavioral strategies estimated in Table 6. It is difficult to discern from reading the three blocks of coefficients in the table which strategy is more responsive to cost conditions, which is more responsive to rivals’ moves, which exerts more pressure on rivals, etc. It is difficult to tell how the behavioral strategies interact with each other. The counterfactuals provide a way of to understand such issues. Our intent is not to forecast market outcomes from a change in the mix of firms’ structural parameters. To be more concrete, consider the counterfactual in which the mix of firms is changed to make them all group 1. We cannot use this counterfactual to forecast what would happen if all firms were given the group-1 structural parameters. The strategy played by group-1 firms is presumably tailored to the observed mix of firms and would change against a different mix of rivals. However, our counterfactual simulations allow us to learn whether the behavioral strategy estimated for group 1

in Table 6 leads to more aggressive pricing than the other groups' strategies. Because it is a small change in market structure, the simulations adding a single firm of a given type raise less of a concern that the market structure change would affect behavioral strategies.

The simulations use the time series of cost shocks observed in the data and allow for a Poisson entry process using the average probability observed over the whole sample.¹⁸ In each simulation, we create a set of twenty-four simulated firms, assigning them starting margins by drawing randomly from the empirical margin distribution in the data. From those margins, given costs, we then can back out prices and generate ranks. We then allow cost to follow the observed series and allow the simulated firms' prices to evolve as governed by the empirical models in Table 6, essentially simulating a year of market activity hour by hour. The mix of different strategic groups can be freely specified in a simulation run. For instance, to simulate the dynamics of price competition with a group of firms entirely composed of active, low-price firms, we specify that all simulated firms' hour-by-hour decisions are based on the group-1 empirical estimates.

We performed seven different simulations. The simulation reported in column (2) reproduces the actual mix of the three strategic groups in the data. The simulations in columns (3)–(5) each add a single firm of a given type (so a 25th firm) to the market. The simulations in columns (6)–(8) change the market structure even more substantially by moving them all into one strategic group. Table 7 shows the gross margins from the simulations and Table 8 the gross profits. Each type of simulation was run 100 times, the median computed for each run, and then these medians were averaged over the 100 runs. We used the median for each run to reduce simulation variability, but results were similar using the average of the means rather than the average of the medians. The standard deviation across the 100 simulation runs are shown in parentheses as a measure of simulation variability.

The simulation maintaining the actual mix and density in column (2) produces similar but slightly smaller margins than in the actual data, reported in column (1). For all firms, margins averaged 1.6% in the data but 0.7% in the simulations on average across all firms. The ordering of margins by group was maintained, with group 1 having the lowest margins, group 2 higher margins, and group 3 the highest. The breakdown of simulated margins by strategic group matches the breakdown in the actual data reasonably well. For example, the simulated margin for group 1 is -1.6 , close to the -1.8 observed in the actual data.

¹⁸The probability of new entry or re-entry after a gap in activity averages 0.28% per hour in our sample. Conditional on entry, the probability that the entrant occupies one of the ranks 1–14 is 1/34 one of ranks 15–24 is 1/17. Entry knocks the highest-priced firm off the list so that the simulated market always maintains 24 firms.

Table 8 converts the results for gross margins in terms of profit per day. See the table notes for the exact formula, which makes some minor adjustments for merchant credit-card fees, input loss, and any wedge between shipping fees and costs. The key adjustment in calculating profit is multiplication by quantity per day, which in turn depends most crucially on the firm's rank. The profit figures should be interpreted with caution because firms earn profit from other sources not captured in the formula. As studied in Ellison and Ellison (2009a), consumers attracted to a firm's website by low offers for RAM modules can be sold premium RAM models or other add-ons with higher margins, and these additional profits are not captured in the formula. These figures are thus better seen as providing a lower bound on profit. The profits in Table 8 mirror the margins in the previous table. In particular, the high-margin strategy pursued by group-3 firms has some rationale. Even if these firms earn no more than this lower bound, they can manage to sell enough to make a profit of \$16.70 per day from this one item. In the simulations, the figure for these firms is lower but still substantial at \$9 per day. The lower bound on profit for the other firms is negative for group 1 and only slightly greater than 0 for group 2, perhaps an indication of the profits they earn from other sources such as add-on sales. The argument that there must be other sources of profit at least making up the difference between the group-1 and group-3 averages is based on revealed preference: a group-1 firm could have adopted the group-3 strategy and earned that average profit but chose not to.

The simulations in columns (3)–(5) add a single firm to the market from each group in succession. The addition of a group-1 firm leads to a reduction in margins for firms in all strategic groups, by 0.4% for group 1, 0.5% for group 2, and 0.6% for group 3. Adding a group-2 firm also reduces margins for all groups, but the effect is smaller than adding a group-1 firm. Margins for all groups fall when a group-3 firm is added, but the effect is even smaller than adding a group-2 firm. Group-1 firms seem to provide the most competitive pressure, followed by group-2 firms, followed by group-3 firms, in order of their average margins. However, the competitive pressure seems to be felt fairly uniformly by rivals regardless of group.

Columns (6)–(8) show how margins and profits change when instead of the actual mix all firms are assumed to use the behavioral strategy of one of the strategic groups. Moving all the firms into group 1 reduces the margins of group-1 firms from -1.6% to -2.8% . Moving all firms into group 2 increases margins for group-2 firms from 5.0% to 5.7% . The reduction in competitive pressure from converting of group-1 to group-2 firms dominates the increase from converting group-3 into group-2 firms. Moving all firms into group 3 increases margins of group-3 firms about a percentage point, from 8.6% to 9.5% . This fairly small

increase indicates that group-3 firms still exert competitive discipline on each other. The effects on profits is more substantial, rising from \$9.0 to \$15.0 per day.

9. Conclusion

Using the results from our estimation of a dynamic pricing model, we can return to the Introduction and address a number of questions raised there. The Introduction asked whether firms in our market targeted ranks or margins with their price changes. The results from Table 3, in particular the significant coefficients on *Margin* and *NumBump*, suggest that both targets typically factored into firms' price changes.

We found ample evidence of heterogeneity in firms' strategies. The wide gap between the price series for the two firms shown Figure 3 provided obvious evidence of price dispersion. More subtle but still apparent was the difference in the timing of the firms' price changes. We provided a number of more formal tests of heterogeneity. A likelihood-ratio test strongly rejected pooling all firms in a single price model. In the interest of parsimony, we used an agglomeration method to group the firms into three strategic groups and re-estimated the pricing model separately for each of those. Figure 7 allowed us to visualize the essential character of each of the three groups: the first group included low-price firms that adjusted prices frequently, the second group moderate-price firms that rarely adjusted prices, and the third group yet higher priced firms that again adjusted prices frequently. The estimates of the dynamic pricing model reported in Table 6 while not differing qualitatively, differed quantitatively across groups.

The simulations in Tables 7 and 8 provided a gauge as to the importance of these quantitative differences in strategy. While adding a firm using the strategies estimated for any of the three groups would intensify competition, the first group exerted the greatest competitive pressure, followed by the second group, followed by the third group. Interestingly, the percentage-point reduction in margins or profits was fairly uniform across strategic groups. So for example, the addition of a group-2 firm reduced the margins of group-1 and group-3 firms by about the same number of percentage points as group-2 firms. Besides gauging the importance of strategy heterogeneity, the simulations also served as a specification check on the dynamic pricing model.

We found substantial pricing inertia in the Pricewatch market during our sample period. In fewer than a third of the days in which the firm faced a cost change or rival price change that bumped the firm from its previous rank did the firm respond with a price change of its own. The sources of pricing inertia postulated

by Davis and Hamilton (2004) do not seem to fit our market. Menu costs were small because changing price was a simple matter of entering a new number in the Pricewatch database. Information lags were also not important—firms could monitor rivals’ prices on the website continuously and received updated cost information when reordering memory modules from the wholesaler—nor was consumer acceptance—since there was little branding and repeat business from consumers. Firms may have tried to maintain a strategic rapprochement among themselves, although this might explain inertia in the face of cost changes, not bumps in ranking.

Instead, our results point to a different source of pricing inertia: managerial inattention. Several pieces of evidence pointed in this direction. The coefficients on *SinceChange* in Table 3 indicate that the threshold for both price increases and decreases expand over time, consistent with a declining hazard of price change and thus price changes clustered in time followed by episodes of inattention. As shown in Table 6, not all strategic groups exhibited this behavior, but it was strongly evident for firms showing the most activity, firms in group 1. Further evidence of managerial inattention came from the coefficients on *Night* and *Weekend* in Table 3. The thresholds for pricing activity expanded greatly during these periods, implying that managers only rarely changed price during these periods, exactly when it was expensive for him or her to do so. More refined evidence came from examining the difference in price activity over the course of the day for suppliers on the East versus West Coasts. Measured in Eastern Time, the day starts three hours later for managers on the West than the East Coast. Using a formal test of first order stochastic dominance, we showed that pricing activity was also significantly later on the West Coast. Interestingly, the modal price change was not simply shifted three hours later: the mode for East Coast suppliers was in the morning in Eastern Time but for West Coast suppliers was eight hours later in the evening Eastern Time.

Prices generally did not cycle in our market. As is typical for computer components, costs had a strong secular downward trend, and prices generally followed suit. However, cycles did show up in ranks, shown most vividly in Figure 7. Cycles were different from the Edgeworth variety, documented by previous work on sponsored-search and online-advertising auctions (Edelman and Ostrovsky 2010, Zhang 2010) as well as a handful of papers on retail gasoline cited in the Introduction. An Edgeworth cycle involves a gradual decline in prices as firms sequentially undercut each other, remaining near cost until one jumps up to the monopoly price, causing others to relent as well. Translated into ranks, the largest and most sudden adjustments associated with Edgeworth cycles would be upward. In our market, firms typically adjusted price

downward, resulting in sudden drops not rises in ranks. The gradual rise in a firm's rank was typically due to pricing inertia in the face of rival price drops.

In on-going work, we are estimating a structural dynamic model that provides estimates of unobserved marginal cost sources (such as the cost of order fulfillment) and revenue sources (such as selling high-margin items to the consumer in addition to or in place of the listed low-margin memory module) as well as costs to the manager of attending to the market each instant and of computing a price change when a threshold is reached.

References

- Aguirregabiria, Victor and Pedro Mira. (2007) "Sequential Estimation of Dynamic Discrete Games," *Econometrica* 75: 1-53.
- 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.
- Barro, Robert. (1972) "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies* 39: 17-26.
- 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.
- Bhaskar, V., S. Machin, and G. Reid. (1991) "Testing a Model of the Kinked Demand Curve," *Journal of Industrial Economics* 39: 241-254.
- Bhaskar, V. and Fernando Vega-Redondo. (2002) "Asynchronous Choice and Markov Equilibria," *Journal of Economic Theory* 103: 334-350.
- Bishop, J., A. J. Formby, and P. Thistle. (1992) "Convergence of the South and Non-South Income Distributions, 1969-1979," *American Economic Review* 82: 262-272.
- 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.
- Cahuc, Pierre and Hubert Kempf. (1999) "Asynchronized Multiperiod Commitments and Cycles," *Journal of Economic Behavior and Organization* 40: 387-407.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. (2004) "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics* 119: 861-898.
- Carleton, Dennis W. (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.
- 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.
- Davidson, Russell and Jean-Yves Duclos. (2000) "Statistical Inference for Stochastic Dominance and for the Measurement of Poverty and Inequality," *Econometrica* 68: 1435-1464.
- Davis, Douglas D. and Oleg Korenok. (2009) "Posted Offer Markets in Near-Continuous Time: An Experimental Investigation," *Economic Inquiry* 47: 449-466.

- 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. (2004) "An Alternating-Move Price-Setting Duopoly Model with Stochastic Costs," *International Journal of Industrial Organization* 22: 997–1015.
- Eckert, Andrew and Douglass West. (2004) "Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations," *Journal of Law and Economics* 47: 245–271.
- Edelman, Benjamin and Michael Ostrovsky. (2010) "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems*, forthcoming.
- Edgeworth, Francis. (1925) "The Pure Theory of Monopoly," in *Papers Relating to Political Economy*, vol. 1. London: MacMillan, 111–142.
- 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.
- Fishman, Arthur. (1992) "Search Technology, Staggered Price-Setting, and Price Dispersion," *American Economic Review* 82: 287–298.
- Golosov, Mikhail and Robert E. Lucas, Jr. (2007) "Menu Costs and Phillips Curve," *Journal of Political Economy* 115: 171–199.
- Greene, William. (1981) "On the Asymptotic Bias of the Ordinary Least Squares Estimator of the Tobit Model," *Journal of Econometrics* 21: 195–212.
- Hansen, Per Svejstrup, *et al.* (1996) "Asymmetric Adjustment in Symmetric Duopoly," *Economic Letters* 53: 183–188.
- Harrington, Joseph E., Jr. (1995) "Experimentation and Learning in a Differentiated-Products Duopoly," *Journal of Economic Theory* 66: 275–288.
- Heckman, James. (1976) "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models," *The Annals of Economic and Social Measurement* 5: 475–492.
- 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 Benjamin A. Malin. (forthcoming) "Microeconomic Evidence on Price-Setting," in Benjamin Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, vol. 3. Amsterdam: North-Holland.
- Lach, Saul. (2002) "Existence and Persistence of Price Dispersion: An Empirical Analysis," *Review of Economics and Statistics* 84: 433–444.

- Lal, Rajiv and Carmen Matutes. (1989) "Price Competition in Multimarket Duopolies," *Rand Journal of Economics* 20: 516–537.
- Lau, Sau-Him Paul. (2001) "Aggregate Pattern of Time-dependent Adjustment Rules, II: Strategic Complementarity and Endogenous Nonsynchronization," *Journal of Economic Theory* 98: 199–231.
- Leufkens, Kasper and Ronald Peeters. (2008) "Intertemporal Price Competition with Exogenous Demand Shocks," *Economic Letters* 99: 301–303.
- 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.
- Maskin, Eric and Jean Tirole. (1988) "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica* 56: 571–599.
- McFadden, Daniel. (1989) "Testing for Stochastic Dominance," in T.B. Formby and T.K. Seo, eds., *Studies in the Economics of Uncertainty*. New York: Springer-Verlag.
- Nakamura, Emi and Jón Steinsson (2008) "Five Facts about Prices: A reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123: 1415–1464.
- 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.
- Romesburg, H. Charles. (2004) *Cluster Analysis for Researchers*. North Carolina: Lulu Press.
- 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.
- Zhang, Xiaoquan. (2010) "Finding Edgeworth Cycles in Online Advertising Auctions," Sloan School of Management working paper.

Table 1: Variables in the dataset

Variable	Definition	Mean	Std. dev.	Min.	Max.	Obs.
Firm-level variables						
<i>Price</i>	Current listed price in dollars	66.92	34.66	21	131	126,514
<i>Margin</i>	Gross margin of price over wholesale cost in percent, $100(p-c)/p$	2.64	9.91	-34.67	37.50	126,514
<i>Rank</i>	Rank of listing in price-sorted order	10.62	6.74	1	24	126,514
<i>NumBump</i>	Net number of ranks bumped since last price change	1.33	3.34	-22	21	126,514
<i>SinceChange</i>	Hours since firm last changed its price	113.99	140.79	1	1,112	126,514
<i>Placement</i>	Placement between adjacent firms in price space	0.58	0.42	0	1	126,514
<i>Density</i>	Measure of density in price space of firms with nearby ranks	0.58	0.39	0	3	126,514
<i>FirstPage</i>	Indicates appearance on first webpage (rank 1-12)	0.62	0.48	0	1	126,514
<i>Tenure</i>	Number of days since entered market	74.03	56.68	0	228.83	46,065
Market-level variables						
<i>CostTrend</i>	Wholesale cost trend over previous two weeks	-0.18	0.70	-2.06	1.53	8,185
<i>CostVol</i>	Wholesale cost volatility over previous two weeks	1.62	1.06	0.00	4.36	8,185
<i>Night</i>	Indicates hour from midnight to 8 a.m. EST	0.33	0.47	0	1	8,218
<i>Weekend</i>	Indicates Saturday or Sunday	0.29	0.45	0	1	8,218
<i>PriceLow</i>	Lowest of currently posted prices	62.63	32.90	21	122	8,218
<i>PriceRange</i>	Range across currently posted prices	12.62	4.36	5	29	8,218
<i>AvgMargin</i>	Average across firms' gross margins in percent	4.23	7.23	-12.17	28.03	8,218
<i>NewFirm</i>	Indicates new entry into ranks 1-24 in past two weeks	0.63	0.48	0	1	8,218

Notes: Descriptive statistics for *Tenure* only for subsample of firms entering after start of sample (the variable is left-censored for other firms), hence the reduction in observations for that variable.

Table 2: Separate regression estimates of baseline model of price change

Variable	Probit on indicator for price decrease (1)	Probit on indicator for price increase (2)	OLS on size of price increase (3)
Constant	-2.2361*** (0.0435)	-2.6078*** (0.0535)	-0.0135*** (0.0021)
<i>Margin</i>	-0.0027 (0.0027)	-0.0240*** (0.0019)	-0.0016*** (0.0002)
<i>NumBump</i>	0.0452*** (0.0084)	-0.0544*** (0.0092)	-0.0063*** (0.0008)
<i>CostTrend</i>	-0.0709** (0.0290)	0.0264 (0.0381)	-0.0033 (0.0022)
<i>CostVol</i>	-0.0083 (0.0077)	0.0710*** (0.0192)	0.0038*** (0.0013)
<i>SinceChange</i>	-0.0013*** (0.0002)	-0.0013*** (0.0003)	
<i>Night</i>	-0.5686*** (0.1109)	-0.7294*** (0.1182)	
<i>Weekend</i>	-0.4051*** (0.0688)	-0.3284*** (0.0713)	
R^2 or Pseudo R^2	0.060	0.118	0.309
Observations	126,514	126,514	1,213

Notes: Each column is a separate regression. Dependent variable in (1) is an indicator for the event $Price_{i,t} < Price_{i,t-1}$, in (2) is an indicator for the event $Price_{i,t} > Price_{i,t-1}$, and in (3) is $\ln Price_{i,t} - \ln Price_{i,t-1}$. Probit coefficient estimates reported in (1) and (2) and ordinary least squares coefficient estimates reported in (3). Column (3) run on subset of observations i,t for which $Price_{i,t} \neq Price_{i,t-1}$. Heteroskedasticity-robust standard errors clustered by firm reported in parentheses. Sample includes 43 firm clusters. Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 3: Joint maximum likelihood estimates of baseline model of price change

Variable	Latent desire to increase price (1)	Lower threshold (for price decrease) (2)	Upper threshold (for price increase) (3)	Desired size of price increase (4)
Constant		-2.2639*** (0.0515)	2.5804*** (0.0436)	0.0009 (0.0019)
<i>Margin</i>	-0.0048** (0.0022)			-0.0009*** (0.0002)
<i>NumBump</i>	-0.0469*** (0.0066)			-0.0031*** (0.0006)
<i>CostTrend</i>	0.0623*** (0.0230)			-0.0054*** (0.0019)
<i>CostVol</i>		0.0196** (0.0082)	-0.0795*** (0.0176)	0.0012 (0.0010)
<i>SinceChange</i>		-0.0015*** (0.0002)	0.0015*** (0.0003)	
<i>Night</i>		-0.5671*** (0.1100)	0.7029*** (0.1164)	
<i>Weekend</i>		-0.4064*** (0.0689)	0.3114*** (0.0651)	

Notes: Maximum likelihood estimates of coefficients from model in equations (1)–(4). Columns (1)–(4) estimated jointly using 126,514 observations. Estimate of standard deviation of error from equation (2) is $\hat{\sigma} = 0.033$ and of the correlation between errors in equations (1) and (2) is $\hat{\rho} = 0.417$. Log-likelihood is $-4,509.29$. Heteroskedasticity-robust standard errors clustered by firm reported in parentheses. Sample includes 43 firm clusters. Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 4: Maximum likelihood estimates of expanded model of price change

Variable	Latent desire to increase price (1)	Lower threshold (for price decrease) (2)	Upper threshold (for price increase) (3)	Desired size of price increase (4)
Constant		-2.3382*** (0.0759)	2.7711*** (0.0778)	-0.0118** (0.0057)
<i>Margin</i>	0.0061 (0.0045)			-0.0013*** (0.0004)
<i>NumBump</i>	-0.0467*** (0.0068)			-0.0031*** (0.0006)
<i>CostTrend</i>	0.0063 (0.0252)			-0.0053*** (0.0019)
(a) <i>Quantity</i>	0.0037*** (0.0007)			0.0001 (0.0001)
(a) <i>Placement</i>	0.0617 (0.0431)			0.0048* (0.0025)
(a) <i>PriceLow</i>	0.0005 (0.0004)			0.0002*** (0.0001)
(a) <i>PriceRange</i>	-0.0039 (0.0034)			-0.0006* (0.0003)
(a) <i>AvgMargin</i>	-0.0163*** (0.0050)			0.0010** (0.0004)
(a) <i>NewFirm</i>	-0.1575*** (0.0307)			0.0046** (0.0018)
(a) <i>Density</i>		-0.0086 (0.0442)	-0.3663*** (0.0633)	-0.0053* (0.0028)
<i>CostVol</i>		0.0249*** (0.0085)	-0.0677*** (0.0180)	0.0018** (0.0008)
<i>SinceChange</i>		-0.0015*** (0.0002)	0.0013*** (0.0003)	
<i>Night</i>		-0.5731*** (0.1088)	0.7103*** (0.1184)	
<i>Weekend</i>		-0.4056*** (0.0685)	0.3277*** (0.0641)	

Notes: See notes from previous table. Variables added to list in that table marked (a). Estimate of standard deviation of the error from equation (2) is $\hat{\sigma} = 0.032$ and of the correlation between errors in equations (1) and (2) is $\hat{\rho} = 0.433$. Log-likelihood is -4,397.11.

Table 5: Selected descriptive statistics by strategic group

Variable	Group 1	Group 2	Group 3
Means			
<i>Margin</i>	-0.32	2.77	11.81
<i>Rank</i>	7.82	12.95	17.83
<i>NumBump</i>	1.01	2.17	1.75
<i>SinceChange</i>	96.80	218.04	96.47
Number of firms	22	8	13
Observations	82,277	18,012	26,225

Table 6: Maximum likelihood estimates of baseline model allowing parameters to vary across strategic groups

Variable	Group 1				Group 2				Group 3			
	Latent desire to increase price	Lower threshold	Upper threshold	Desired size of price increase	Latent desire to increase price	Lower threshold	Upper threshold	Desired size of price increase	Latent desire to increase price	Lower threshold	Upper threshold	Desired size of price increase
Constant		-2.1981*** (0.0518)	2.5635*** (0.0465)	0.0015** (0.0008)		-2.7583*** (0.1018)	2.4906*** (0.2771)	-0.0186* (0.0096)		-2.4737*** (0.1823)	2.7834*** (0.1330)	0.0046 (0.0082)
Margin	-0.0091*** (0.0028)			-0.0012*** (0.0002)	-0.0050** (0.0021)			-0.0018 (0.0013)	-0.0105** (0.0043)			-0.0007* (0.0004)
NumBump	-0.0557*** (0.0099)			-0.0032*** (0.0008)	-0.0397*** (0.0147)			-0.0038*** (0.0010)	-0.0418*** (0.0130)			-0.0034*** (0.0011)
CostTrend	0.0476* (0.0265)			-0.0052** (0.0023)	0.1275** (0.0548)			-0.0089 (0.0073)	0.0081 (0.0295)			-0.0045** (0.0020)
CostVol		0.0259** (0.0094)	-0.0894*** (0.0201)	0.0002 (0.0012)		0.0158 (0.0429)	-0.1039** (0.0435)	0.0111** (0.0045)		0.0492*** (0.0177)	-0.0965* (0.0492)	0.0021 (0.0028)
SinceChange		-0.0015*** (0.0003)	0.0015*** (0.0003)			-0.0006 (0.0004)	0.0012 (0.0011)			-0.0019*** (0.0005)	0.0002 (0.0005)	
Night		-0.5256*** (0.1205)	0.6837*** (0.1282)			-0.6121*** (0.2110)	0.6624** (0.2843)			-1.0490*** (0.3170)	3.3478*** (0.0919)	
Weekend		-0.3910*** (0.0762)	0.2794*** (0.0717)			-0.2414 (0.2567)	0.5792* (0.3011)			-0.5739*** (0.1801)	0.6311** (0.1721)	
$\hat{\sigma}$		0.033				0.033				0.030		
$\hat{\rho}$		0.430				0.297				0.412		

Notes: Maximum likelihood estimates of coefficients from the same model estimated in Table 3. Firms grouped using clustering procedure described in text. Parameters and standard errors estimated conditional on resulting grouping. Parameter $\hat{\sigma}$ is an estimate of the standard deviation of the error from equation (2) and $\hat{\rho}$ of the correlation between the errors in equations (1) and (2). Heteroskedasticity-robust standard errors clustered by firm reported in parentheses. Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 7: Gross margins from actual data vs. simulations

Firms	Computed from data	Simulated scenario						
		Maintaining actual mix and density	Adding single firm from			Moving all firms to		
			Group 1	Group 2	Group 3	Group 1	Group 2	Group 3
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
All	1.6	0.7 (0.8)	0.3 (0.7)	0.7 (0.8)	0.9 (0.7)	-2.8 (0.5)	5.7 (2.7)	9.5 (1.4)
Group 1	-1.8	-1.6 (0.6)	-2.0 (0.5)	-1.9 (0.6)	-1.9 (0.8)	-2.8 (0.5)		
Group 2	1.7	5.0 (2.2)	4.5 (2.3)	4.7 (2.3)	4.9 (2.5)		5.7 (2.7)	
Group 3	12.1	8.6 (2.5)	8.0 (1.9)	8.2 (2.1)	8.3 (2.0)			9.5 (1.4)

Notes: Figures in column (1) are medians over firm×period observations in the actual sample. To derive the figures in columns (2)-(6), medians over firm×period observations are computed for individual simulations; these are then averaged over the 100 simulations run for each scenario. Standard deviation of the medians across the 100 simulations (thus only reflecting simulation variability) reported in parentheses. All figures are percentages. The simulation maintaining the actual mix puts 15 of the 24 firms in group 1, 4 of the 24 in group 2, and 5 of the 24 in group 3. This division mimics the proportion of observations in our sample associated with each group.

Table 8: Profits from actual data vs. simulations

Firms	Computed from data	Simulated scenario						
		Maintaining actual mix and density	Adding single firm from			Moving all firms to		
			Group 1	Group 2	Group 3	Group 1	Group 2	Group 3
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
All	2.8	2.1 (1.0)	1.6 (0.9)	2.1 (1.0)	2.2 (0.9)	-2.7 (0.7)	9.1 (3.5)	15.0 (1.9)
Group 1	-1.1	-1.4 (1.1)	-1.6 (0.9)	-1.4 (1.0)	-1.5 (1.0)	-2.7 (0.7)		
Group 2	1.5	6.9 (2.6)	6.1 (2.3)	6.4 (2.4)	6.8 (2.3)		9.1 (3.5)	
Group 3	16.7	9.0 (1.7)	8.1 (1.4)	8.4 (1.6)	8.6 (1.4)			15.0 (1.9)

Notes: Profit computed according to the formula $(0.975p - 1.01c + 2.25)q$, where price p is scaled down to reflect the 2.5% credit-card fee charged to merchants, wholesale acquisition cost c is scaled up to reflect loss of units due to theft and breakage, and \$2.25 is added per unit to reflect the excess of shipping and handling charges to consumers (modal value of \$10.75) over supplier's actual shipping and handling costs (\$8.50 according to our best estimates). Quantity q computed from the demand estimates of Ellison and Ellison (2009a). Figures in column (1) are medians over firm \times period observations in the actual sample. To derive the figures in columns (2)-(5), medians over firm \times period observations are computed for individual simulations; these are then averaged over the 100 simulations run for each scenario. Standard deviation of the medians across the 100 simulations (thus only reflecting simulation variability) reported in parentheses. All figures are in dollars per day. The simulation maintaining the actual mix puts 15 of the 24 firms in group 1, 4 of the 24 in group 2, and 5 of the 24 in group 3. This division mimics the proportion of observations in our sample associated with each group.

Figure 1: Example Pricewatch webpage


PRICE WATCH® Query = System Memory PC100 128MB

[Back To - New Components Home](#)

[Back To - "Not Exactly New" Home](#)


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	pboost.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

Figure 2: Effect of rank on demand



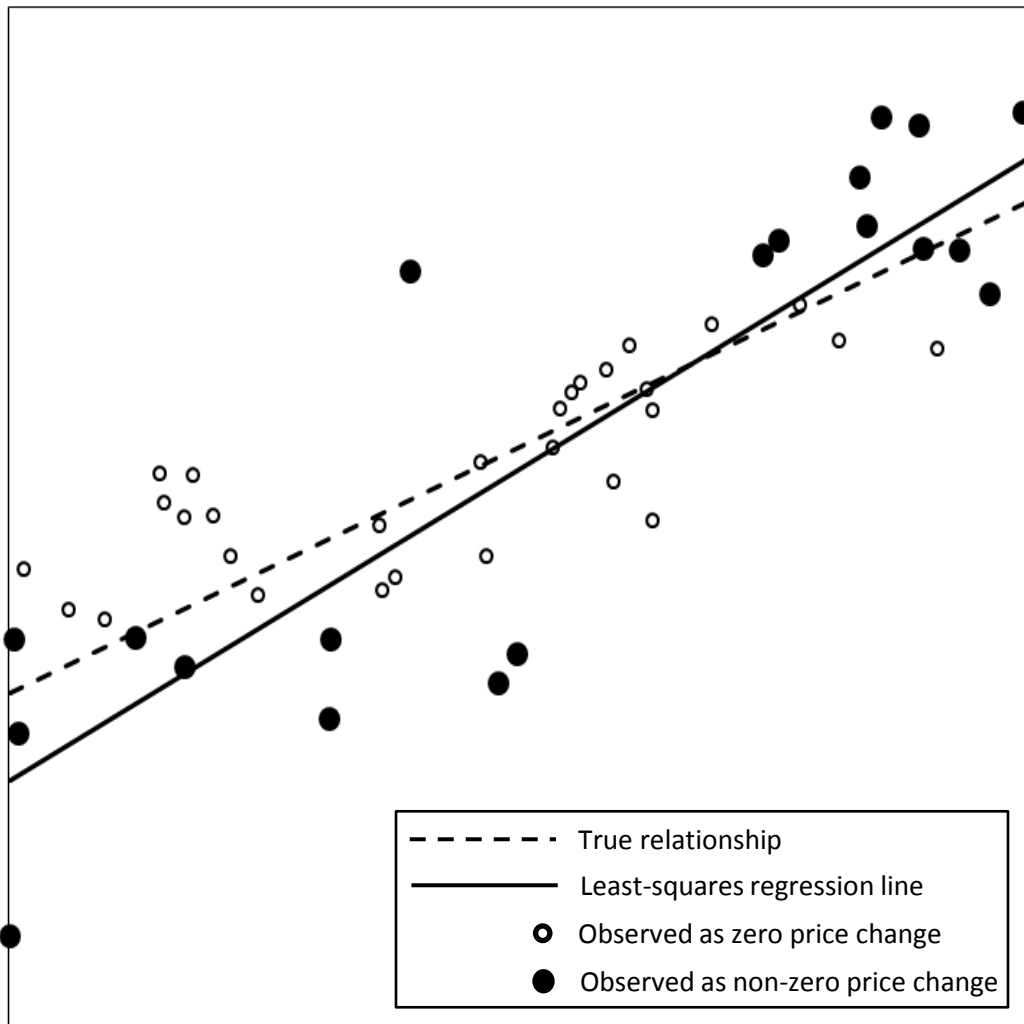
Notes: Estimates of effect of rank on demand from Ellison and Ellison (2009a) using two different specifications. Solid line is from a regression of the natural log of quantity on linear rank. Dots are coefficients from a more flexible specification with rank fixed effects; vertical bars provide 95% confidence interval around fixed effects.

Figure 3: Gross margins for example firms



Notes: Gross margins for two example firms selected on the basis of average rank. The example low-ranked firm had the lowest average rank among firms in the sample. The example median-ranked firm had an average rank closest to the middle (12.5) of the range (1-24) of ranks. Last third of the sample period not shown because one of the firms exited the market. Graphs show gaps when firms fell out of the first two Pricewatch pages.

Figure 4: Example of selection bias arising in the price-setting model



Notes: Diagram meant for schematic purposes but based on an actual Monte Carlo simulation with 50 observations. Variable on the horizontal axis is uniform [0,50] random deviate. Price change variable on the vertical axis is the sum of this and a random normal error with mean 0 and variance 100. Assumption is that price change only observed if it is below 4 or above 39; these observations correspond to the filled circles. Dotted line has the true slope of 1; solid line is the regression line estimated on just the filled circles.

Figure 5: Schematic diagram of joint maximum likelihood model


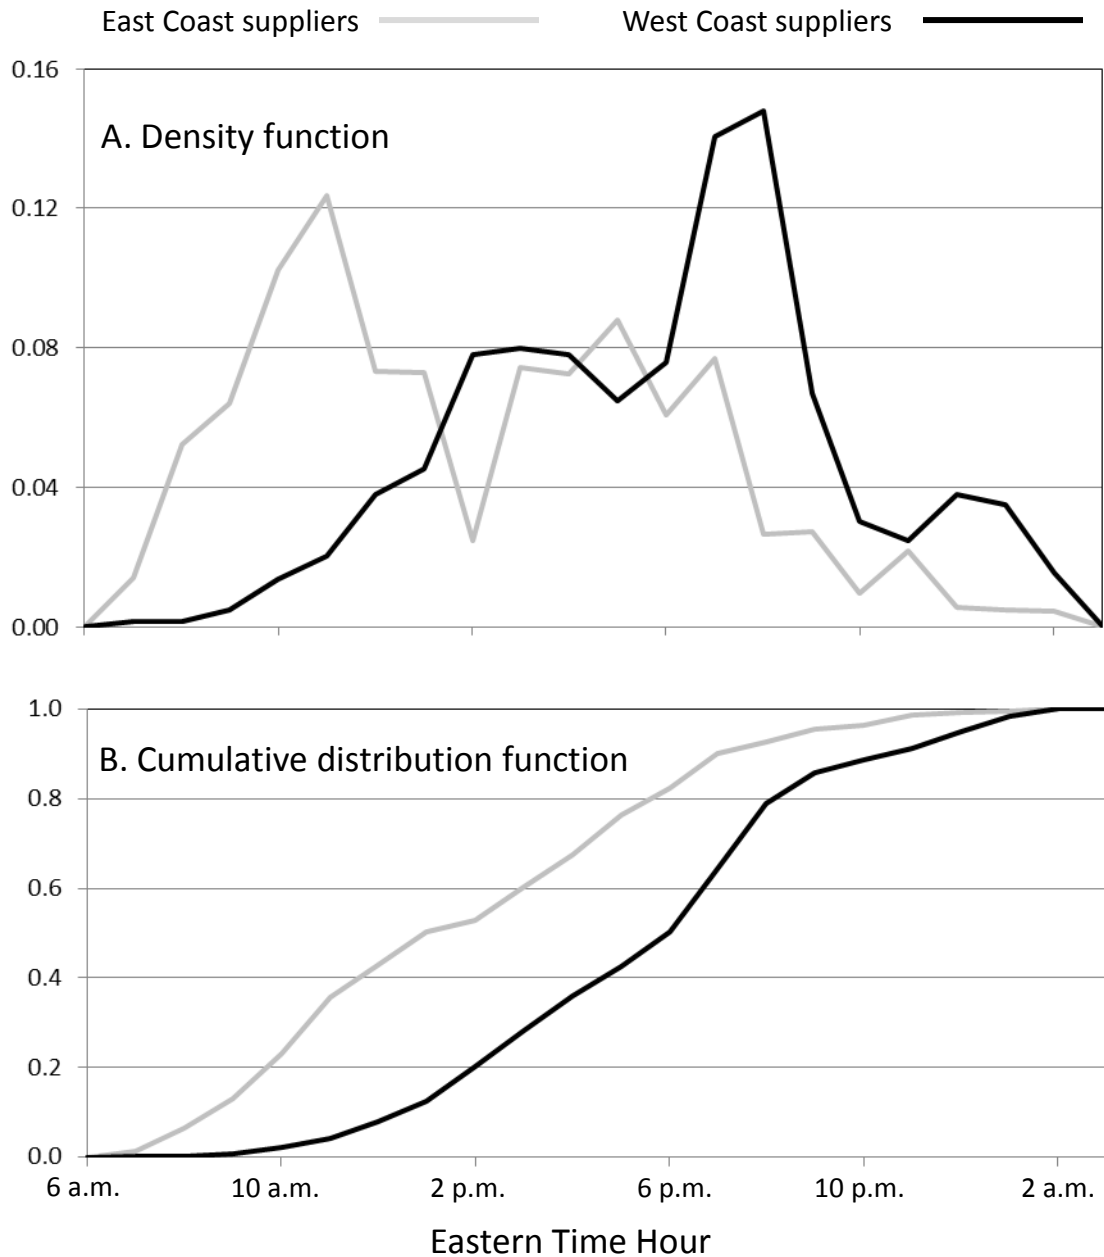
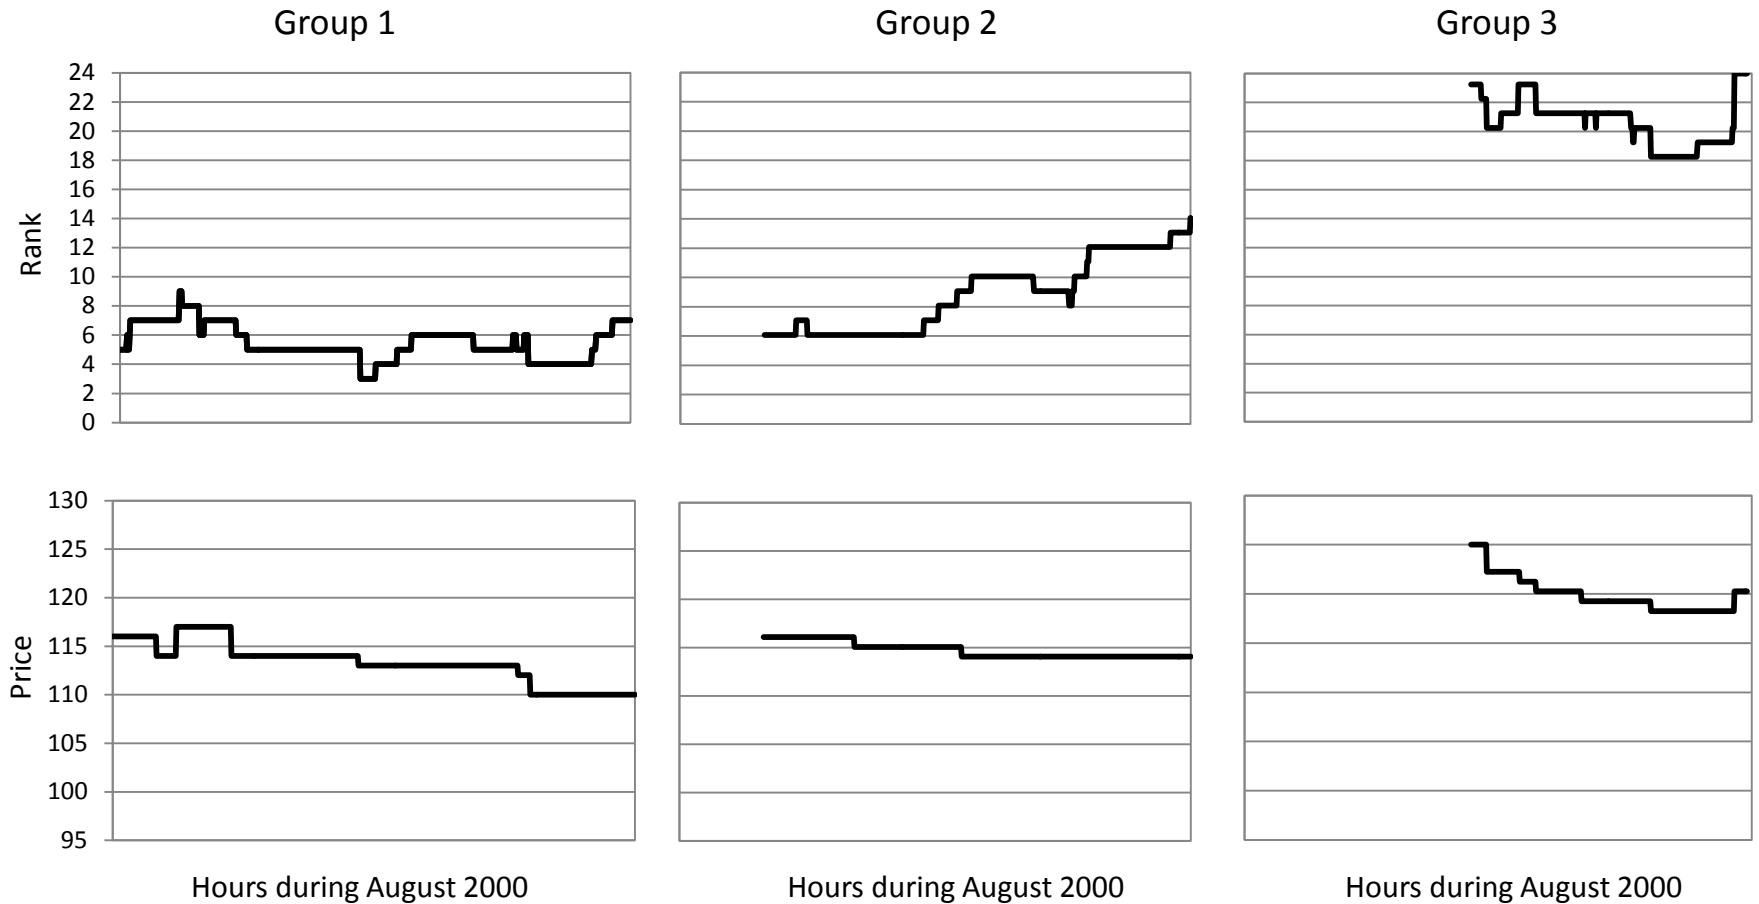


Figure 6: Conditional probability of price change in different time zones over day



Notes: Calculations based on joint maximum likelihood model 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. Panel A computed by taking the covariates for each observation in the dataset, resetting the indicators for supplier location and Eastern Time Zone hour so that they are 0 for all but the labeled case, computing the predicted probability θ_{lit} of a price change by a supplier in location $l \in \{E, W\}$ in hour h for each observation it using model coefficients, and averaging over dataset $\theta_{lh} = \sum_{i=1}^N \sum_{t=1}^T \theta_{lit} / NT$. Unconditional probabilities θ_{lh} converted into conditional probabilities in two steps. First the odds ratio $o_{lh} = \theta_{lh} / (1 - \theta_{lh})$ is computed; second the ratio is normalized $\tilde{o}_{lh} = o_{lh} / \sum_{\tau=6 \text{ a.m.}}^{2 \text{ a.m.}} o_{l\tau}$ to form a proper density function (i.e., integrates over its range to hours to 1). Panel A graphs \tilde{o}_{lh} , which can be shown to equal the probability that a supplier in l changes price in hour h conditional on it changing price exactly once during that day. Because the variance-covariance matrix for the probabilities is nearly singular, tests in the text carried out by collapsing single hours into two-hour blocks; the associated graphs are similar. Hours from 3 a.m. to 6 a.m. omitted because few observed price changes then and none in at least one of the locations.

Figure 7: Price and rank series for representative firm from each strategic group



Notes: Series start later for firms in groups 2 and 3 because they entered during the month.