Search and Wholesale Price Discrimination

Guillermo Marshall*
g-marshall@u.northwestern.edu

December 23, 2013

Job market paper

Abstract

Many markets for homogeneous goods feature market power and heterogeneity in the prices paid by buyers. Search costs are a common explanation for this phenomenon and are a concern as they generate inefficiencies. In this paper, I study a competitive market for homogeneous goods and, by exploiting a unique dataset, I find three facts that are opposite to what one would expect from a market with these characteristics. First, sellers enjoy market power. Second, one can find customers paying 50 or 60% more than others for the same product at the same day. Third, price differences are systematic at the buyer level, providing evidence that sellers actively practice price discrimination. Inspired by these facts and by evidence supporting search costs as the source of market power, I propose and estimate a structural search model for two purposes. First, to measure how the market power generated by search costs affects welfare and, second, to study how price discrimination may magnify or reduce the welfare effects of search costs by altering competition intensity. My results address two important issues. First, search costs imply price distortions that generate a loss in total surplus that is about two-thirds of the welfare loss when shifting from perfect competition to monopoly. That is, even for a competitive market for homogeneous goods, search costs can have a severe effect on welfare. Second, price discrimination increases total surplus by as much as six percent relative to when sellers set uniform prices. The increase in welfare can be partially explained by price discrimination increasing search incentives and, hence, intensifying competition.

*Department of Economics, Northwestern University; e-mail: g-marshall@u.northwestern.edu. Acknowledgements: I am especially grateful to Igal Hendel for his guidance and support. I am also grateful to Aviv Nevo and Rob Porter for their help and support. I also thank Germán Bet, Laura Doval, José Espín, Aanchal Jain, Chris Lau, Fernando Luco, Álvaro Parra, Esteban Petruzzello, Tiago Pires, Anthony Wray, and Jaber Zarezadeh for helpful suggestions and conversations, and seminar participants at Northwestern University and Pontificia Universidad Católica de Chile. All mistakes are my own.

1 Introduction

In contrast to the Bertrand paradox, many markets for homogeneous goods feature market power and heterogeneity in the prices paid by buyers. Examples include the online markets for books or DVDs and wholesale markets for commodity goods. Many economists have studied how price dispersion can be sustained in equilibrium, and have proposed the search costs faced by buyers as one explanation. Search costs are a concern because they reduce efficiency by limiting customer mobility to low-price sellers. However, search costs are not the only factor that can cause price heterogeneity and that affects efficiency. In markets where sellers choose not to commit to posted prices, price dispersion across sellers may coexist with price discrimination.

I study the market for wholesale food in an urban area in the US, which is an example of a competitive wholesale market for homogeneous goods. The sellers are food distributors, and the buyers are restaurants that make repeated purchases. By exploiting a unique dataset with data at the transaction level, I find three facts that are opposite to what one would expect from a competitive market for homogeneous goods. First, sellers enjoy market power. I will discuss below evidence supporting search costs as the source of market power. Second, there is a significant extent of price and markup dispersion for each product. To be more specific, one can find customers paying 50 or 60% more than others for the same product at the same day. Third, these differences in prices and markups are systematic at the buyer level, providing evidence that sellers actively practice price discrimination.

Inspired by these facts, I propose and estimate a structural search model to measure the welfare implications of the market power generated by search costs. In the model, search costs make it costly for buyers to move to low-price sellers, which gives sellers market power. By not committing to posted prices, sellers set buyer-specific prices by trading off static profits with changes in the expected value of serving a customer.

The market power generated by search costs affects welfare for two reasons. First, sellers raise prices above marginal cost, which has an unambiguous negative effect on total welfare as the quantity purchased by buyers is price sensitive. Second, customer heterogeneity gives sellers incentives to price discriminate, which has an ambiguous effect on both total welfare and profits. In addition to the conflicting welfare effects of price discrimination in the case of a monopolist, price discrimination may increase the propensity to search and, hence, intensify competition.[1] Although it is always a dominant strategy for a seller to price discriminate, intensified competition can make equilibrium profits decrease with price discrimination.

[1] In the case of a monopoly seller, the effect of price discrimination on welfare depends on how the distributional inefficiencies associated with serving customers at different marginal valuations compare to the potential gains in customer surplus if there is an increase in total quantity (Schmalansee 1981, Varian 1985). The effect of price discrimination on competition is discussed in Holmes (1989).

The objective of this paper is twofold. First, I measure the market power generated by search costs and quantify the implications of market power for welfare. Second, I answer the empirical question of how price discrimination may magnify or reduce the welfare implications of search costs. This analysis sheds light on how search costs affect the performance of markets and how price discrimination may change competition intensity by altering search incentives.

I find important results in two dimensions. First, the market power generated by search costs implies price distortions that cause a reduction in total surplus of about two-thirds of the loss in total surplus when shifting from perfect competition to monopoly. This result informs us that even for competitive markets for homogeneous goods, the welfare effects of search costs can be quite severe. Second, price discrimination increases total surplus by as much as 6 percent relative to the case when sellers set uniform prices. Although the welfare effect is modest, price discrimination helps to reverse some of the negative consequences on welfare by increasing search incentives and, hence, intensifying competition. Intensified competition offsets gains from price tailoring and can make the equilibrium profits decrease by more than 20 percent.

The strategy I follow to achieve my objectives combines descriptive exercises and a structural analysis. First, I use descriptive exercises to provide evidence supporting the existence of search costs and to show how sellers make use of the market power generated by search costs. Second, I propose and estimate a search model to study the welfare consequences of search costs and price discrimination.

The data I use in this study contain detailed information on all transactions that were completed in an eight month period by a single food distributor. Each record is a transaction, defined as a customer–product–time period combination, and includes information such as the unit price, unit wholesale cost, quantity, and customer observables. The richness of this data is crucial for my study. The panel structure of the data allows me to observe interruptions in business relationships, which let me understand how and when customers decide to move their business to an alternative seller. Moreover, the wholesale cost information associated with each transaction allows me to compute markups and understand how prices depend on customer heterogeneity.

I find the following evidence consistent with search costs. First, the probability of a customer–seller breakup is positively correlated with the quoted price. This relation between breakups and prices is consistent with the comparative static result of a sequential search model in which the higher the price, the more likely it is that the price will surpass the customer's reservation price and trigger search. Second, buyers rarely sever business relationships when faced with a price decrease, but that they do so when faced with positive price changes. This is also consistent with buyers following the reservation price rule. A fixed reservation price implies that only a price increase can trigger search, as buyers will never reject prices that are lower than prices that they have already accepted (and, as a consequence, are below their reservation prices). Third, business relationships usually last several periods. This inertia in buyer–seller relationships is also consistent with a search model, since as long as the seller does not give the buyer incentives to search, the buyer will find it optimal to continue transacting with the same seller.

In relation to price dispersion and price discrimination, I find that markups vary substantially for each product, and that on average two-thirds of the product-specific markup dispersion is between-customer variation, which stresses the importance of customer heterogeneity for understanding prices. In particular, by using two proxies for search cost, I find evidence that sellers set higher markups for customers who face higher search costs, which is consistent with the pricing in a search model where sellers trade-off static profits with the probability of serving a customer.

In the model, customers are heterogeneous and characterized by a customer type. Customer types differ both in their demand function and in the distribution from which they independently draw search costs in every period. Customers search for sellers sequentially, and the equilibrium search probability is increasing in price with a price elasticity that varies by customer type. Sellers offer a single product and face wholesale costs that are independently drawn from a cost distribution that is uniform to all sellers. Price dispersion comes from both cost heterogeneity across sellers and sellers practicing price discrimination. Inertia in buyer–seller matches is a consequence of both search costs and the assumption that customers get a free quote from the seller with whom they had most recently transacted. Customers making repeated purchases adds dynamics to the sellers' pricing problem, as by serving a customer a seller realizes static profits and also gains the option to serve the customer once again in the next period. I assume sellers know customer types, and as a consequence, sellers set buyer-specific prices by trading off static profits with the change in the expected value of serving a customer.

In my framework, as opposed to in the literature of the estimation of search costs, the argument of identification of search costs is constructed for when one observes transaction level data.[2] The identification of search costs will come from the observed prices balancing the tradeoff between the gains of increasing profits and the losses of lowering the likelihood of serving the buyer, which depends on the search probability. Since the search probability depends on the distribution of search costs, I exploit the optimality of the observed prices to identify the distribution of search costs. I also face the identification difficulty that while the questions of this paper involve an industry-wide analysis, I only have data from a single food distributor. To overcome this difficulty, I assume that the data were generated by a symmetric equilibrium. Under the symmetry assumption, it suffices to observe a single firm following the equilibrium strategy to identify the primitives of the model.

Using the estimates of the structural model, I perform two counterfactual exercises. First, I measure both market power and the welfare loss generated by search costs. I implement this analysis by comparing the current market structure to perfect competition and the monopoly case. Second, I measure how price discrimination affects competition and welfare, relative to the case where sellers set uniform prices.

[2]See Hortacsu and Syverson (2004) and Hong and Shum (2006) for papers on the estimation of search costs.

First, my results indicate that the market power generated by search costs implies a loss in total surplus of about two-thirds of the loss in welfare when shifting from perfect competition to monopoly. This result is informative about how much attention we should pay to search costs as search costs can severely harm welfare even when considering competitive markets for homogeneous goods. Second, price discrimination increases total surplus by as much as 6 percent relative to the case when sellers set uniform prices, and that a key effect behind the increase in welfare is that price discrimination intensifies competition on average. The effect of price discrimination on competition is given by how the customer type–dependent equilibrium price distributions change when shifting from sellers setting uniform prices to sellers practicing price discrimination. As a consequence of the type–dependent equilibrium price distributions changing, the expected gains of requesting an extra price quote also change for the different buyers. A higher expected search reward increases the propensity to search and, hence, intensifies the competition for preserving customers as sellers are forced to lower their prices in order to avoid losing business. Customers benefit from increased competition, but sellers can see the equilibrium profits of serving a given customer decrease by as much as 20 percent, as increased competition can more than offset the gains of tailoring prices.

1.1 Literature review

This paper is related to several strands in the literature. A large body of theoretical work has studied the welfare effects of price discrimination. Schmalansee (1981) and Varian (1985) study the case of a monopolist, in which price discrimination generates both distributional inefficiencies and a change in total output. Katz (1984), Katz (1987), Holmes (1989), Corts (1998), Yoshida (2000), and Armstrong and Vickers (2001) study the case of competing firms practicing price discrimination, which incorporates into the analysis the fact that price discrimination has an effect on how intensely firms compete.

A few empirical papers have recently studied price discrimination in the context of vertical or wholesale markets. Dafny (2010) takes a reduced-form approach to provide evidence of both market power and price tailoring in the market that gathers health insurance carriers and employers. Villas-Boas (2009) and Grennan (2013) focus on studying the welfare effects of price discrimination in the wholesale market for coffee and in the market for a medical device, respectively. In Villas-Boas (2009), firms derive market power by product differentiation and display systematic cost heterogeneity. Grennan (2013) models bargaining between upstream and downstream firms, which affects the results of the welfare analysis as the effect of price discrimination on competition has an interaction effect with bargaining ability.

Other empirical papers studying how price discrimination affects welfare in diverse economic environments include work by Leslie (2004), Hendel and Nevo (2011), and Marshall (2013).

This paper also relates to the literature studying search costs. On the theoretical side of the literature, a number of papers have studied the effect of search costs on market power and on understanding the effect of search costs on the nature of competition between firms. Diamond (1971) finds, in the context of a simple model, that buyers facing search costs leads to the monopoly outcome. Several papers have extended the model by Diamond (1971) and have found that an equilibrium with price dispersion can exist if the model features asymmetries between sellers or buyers or both. For instance, a non degenerate equilibrium price distribution is characterized in a model with wholesale cost heterogeneity across sellers (Reinganum, 1979) and also in a model with customer heterogeneity in search costs (Varian, 1980).

On the empirical side of the search literature, Sorensen (2000), Hortacsu and Syverson (2004) and Hong and Shum (2006) document patterns of price dispersion in different markets and provide evidence that search costs explain the patterns of price dispersion that are found. Hortacsu and Syverson (2004) and Hong and Shum (2006) also provide methodologies to estimate search costs using the observed distribution of prices. Hong and Shum (2006) present a framework to estimate search costs using price data only. The authors only allow for heterogeneity in search costs and impose mixed strategy equilibrium conditions to the observed price distribution to recover the distribution of search costs. Hortacsu and Syverson (2004) extend the approach in Hong and Shum (2006) by allowing for heterogeneity in marginal costs and product differentiation. The authors use market shares and prices, impose an equilibrium condition to the observed prices, and exploit the optimality of prices and the demand functions predicted by the model to recover the primitives. The approach in both Hong and Shum (2006) and Hortacsu and Syverson (2004) differs from the one in this paper in the sense that their papers use "aggregate data" while I use transaction level data. My approach does not require imposing an equilibrium condition to the observed prices, instead, it only requires imposing an optimality condition to each observed price. When one only observes prices from completed transactions and there is price discrimination, it is more appropriate to only impose the optimality of each observed price as opposed to assume that the (censored) distribution of prices is the equilibrium price distribution.

2 Data and descriptive analysis

2.1 Background and data description

The data I use in this study were collected from a food distributor located in an urban area in the US. The distributor serves roughly 1,000 customers, which are mainly restaurants, and sells several hundred different products. The sample period comprises eight consecutive months.

The distributor operates with a number of salespeople who are the intermediaries between the distributor and the customers. Each salesperson manages a number of accounts and decides the prices for each transaction involving their accounts. The sales force is not constrained regarding pricing decisions, but it receives profit-maximizing incentives. The salespeople keep a fixed share of the profits of each transaction they make, and if an item is sold below cost, the salesperson has to reimburse the distributor for the losses. Salespeople quote unit prices, and only if the customer places an order, the salesperson learns how much the customers want to purchase. The distributor restricts its operation to linear prices.

In the data, each record is a transaction (defined as a customer–product–time combination). For each transaction, the data includes the date, identifiers for the customer and product, the unit price, the unit wholesale cost, the quantity transacted, and an identifier for the salesperson that completed the transaction. In addition, the data also includes the name, and zip code of each customer.

A salient feature of prices is captured in Figure 1. This figure shows the variation in markups across transactions for two products, and summarizes for a larger number of products, how the product-specific markup variation can be decomposed in variation between customers (i.e., cross section) and variation over time. From the graphs one can observe that there is an important extent of dispersion for both products. The markups for Product B are more disperse (measured in terms of the range of the data) relative to Product A, and the dispersion takes different shapes for both products. About 44 percent of the overall variation in markups of Product A consists of between-customer variation, while for Product B, it is about 74 percent. The table in Figure 1 shows summary statistics of how much of the product-specific markup dispersion can be explained by between-customer variation for a wider range of products. As can be seen, on average, almost two-thirds of the product-specific markup variation is systematic at the customer level. This dispersion and how it relates to customer heterogeneity provide clear evidence of upstream firms enjoying market power and actively practicing price discrimination.

In the next subsections, I discuss in detail the behavior of buyers and sellers, respectively.

2.2 Customer behavior

In this subsection, I study how buyers decide to break business relationships. For the purposes of this analysis, I consider a breakup as a customer-product specific event, and I define the measure as a "longer than usual" time gap between purchases.[3] Having defined a breakup, it is important to note that customers may have different ways of behaving when breaking relationships. One type of customer may terminate a relationship immediately upon observing a price that is "too high", a second type may break the relationship but still make a last dissatisfied purchase, while a third

[3] A relationship is classified as broken if the time from the last recorded purchase to the next purchase (if any) is greater than the maximum between 50 business days and the maximum time between purchases. In most cases, however, the time gaps after a breakup extend beyond the end of the sample period, meaning that business is never observed to resume. By definition, transactions that are close enough to the end of the sample period are never classified as breakups. The results presented here are robust to varying the criterion.

type may break the relationship after a series of relationship-deteriorating price changes. Since the data only includes completed transactions, one has to interpret the data with caution.

Figure 2 shows evidence consistent with customers responding negatively to higher price changes (top-left), and to higher prices (bottom). In the figure, price changes are measured as the difference between the last two prices that the customer was observed to pay, and prices are measured relative to the product-specific average price. That breakups are positively correlated with prices (bottom graph) is consistent with a search model. The higher the quoted price, the more attractive the option of searching for a lower price.[4] That high prices are not rejected with probability one suggests the existence of a search friction, which can partially explain the markup dispersion discussed above.[5]

The finding that customers respond relatively more to a price increase (top-left graph) is consistent with a search model as well. To see this, think of a model where the customers' search costs and the equilibrium price distribution are constant over time, which also makes the customers' reserve prices be constant over time. This, in turn, implies that a customer would break a relationship only followed by a price increase (i.e., the customer would never reject a price that is lower than something she already accepted). This argument, of course, does not explain the breakups followed by a price decrease that are observed in the graph. But before further analyzing this finding, it is important to remark that these two figures (top-left and bottom) do not distinguish between the customer types mentioned above, as I am only using data of completed transactions. That is, they implicitly assume that all customers are of the second type (i.e., make last dissatisfied purchase, and then leave). This raises the question: Were the breakups that were observed together with negative price changes (top-left) due to those negative price changes or due to price changes in an unobserved transaction?

In an effort of finding evidence of the existence of these different customer types regarding their breakup behavior, I perform the following exercise. For each transaction in the data, I impute a counterfactual transaction immediately following each recorded transaction. The counterfactual transactions are placed in time using the date of the recorded transaction, and using the average time between purchases of each customer-product combination. I impute the wholesale cost in the counterfactual transactions ($c_{t+1}^{\text{counterfactual}}$, henceforth) using the wholesale cost data of the actual transactions that were completed that day.[6] Using the wholesale of each recorded transaction (c_t, henceforth) and the imputed wholesale cost of the counterfactual transactions, I compute the

[4] This relationship also holds when using alternative measures of price that are comparable across products (i.e., markup or percentile rank).

[5] The search friction is given by the resources involved in finding a new provider of the good. While in principle there could also be a switching cost, I am not able to separately identify it from the search cost per quote, without having data on the customers' search behavior after breaking a relationship.

[6] I had to exclude transactions that were close enough to the end of sample because I do not observe the evolution of the wholesale cost beyond that date, hence, I cannot impute $c_{t+1}^{\text{counterfactual}}$ for those.

proportional cost change at the counterfactual transaction, given by

$$(c_{t+1}^{\text{counterfactual}} - c_t)/c_t.$$

In Table 1, I compare the counterfactual cost changes by breakup status at t (i.e., whether or not there was a breakup after the recorded transaction at time t), and also by price change at time t. From the table one can see that those that both broke up and experienced a negative price change at t (i.e., before leaving) have a statistically significantly higher average counterfactual cost change than those who faced a negative price change at t but that did not breakup. Due to the higher average cost change, those who were classified as having broken the relationship following a negative price change are expected to have faced a more positive price change than those that faced a negative price change but did not break the relationship. This evidence is consistent with the claim that the breakups following negative price changes were due to (unobserved) price increases in counterfactual transactions rather than being due to negative price change at t. This evidence supports the existence of customers of the first type listed above (i.e., leave immediately upon a breakup).

Going back to Figure 2, and the discussion of how the price change affects breakups, I adjust the data to accommodate the existence of these two different types of customers regarding their breakup behavior. In the top-right graph of Figure 2, I impute the counterfactual cost change as the price change for all customer-products that were observed to break the relationship followed by a negative price change. As can be seen in this modified graph, the findings are consistent with the customers responding to positive price changes, but not doing so to negative price changes. This makes the data consistent with a search model where customers have search costs that are correlated over time.

Figure 3 presents this same data in an alternative way. The figure shows the unconditional cumulative distribution function of both price changes and prices (relative to the product-specific average price) for both breakup statuses, as well as the CDF of price changes when conditioning on the event of a non-zero price change. In all cases, the CDF of the transactions that resulted in a breakup first-order stochastically dominate the CDF of the transactions that did not result in a breakup. The conclusion is similar, breakups were more often observed together with both positive price changes and higher prices, relative to non-breakups.

Table 2 shows the asymmetric response to price changes that was discussed above in a different way. The table presents results of regressions of the breakup variable on prices (relative to the product-specific average price), and on the relative price interacted with a dummy variable of whether or not the customer faced a positive price change. As can be observed, the positive coefficient on relative price is intensified when the observation includes a positive price change. The result is robust to the inclusion of customer and product fixed effects. Note, however, that given the discussion above of the breakups that are attributed to negative price changes, these results are biased towards zero.

While I discussed and presented evidence of the existence of the first two types of customers (i.e., immediate breakup, and breakup making last transaction), in Table 3, I explore the existence of the third type (i.e., breakup caused by a deteriorating relationship). For this, I regress the breakup variable on the price change, and on up to two lags of this variable. I find that the coefficient on the lags are not statistically different to zero, giving no support to the idea of the existence of this third customer type that terminates a relationship after a series of price increments.

To study whether a breakup in one product triggers a customer to breakup in other products, in Table 4, I tabulate the number of breakups in each order placed by a buyer (i.e., for each customer-time combination). A customer-time combination can include purchases in multiple products. As the table shows, most of the customer-time combinations in which a breakup was observed only involved a single breakup. This finding is consistent with externalities in breakup decisions not playing a significant role.

Finally, Table 5 shows the variation of the number of unique products that customers purchase every month, where an observation is a customer-month combination. While one can be worried that some breakups may be due to seasonality of purchases, the results in the table make clear that seasonality –captured at the business type level (e.g. pizza restaurant)– plays a small role in explaining fluctuations of the number of items purchased. This finding, together with the fact that most of the systematic variation of this variable is customer-specific, supports the hypothesis of search behavior.

Summarizing, in this subsection I have shown that the data is consistent with the existence of a search friction. Although customers are found to be somewhat insensitive to prices when deciding to break a business relationship, breakups are positively correlated with price, as predicted by a search model.

2.3 Seller behavior

2.3.1 Markups and customer heterogeneity

In this subsection I study how sellers exercise market power. Table 6 shows the decomposition of the systematic variation in markups, and just as in Figure 1, the results in the table remark the importance of customer heterogeneity for understanding prices. The table shows that only half of the systematic variation in markups is product-specific, and that other observables like the type of restaurant (e.g. pizza or seafood), the salesperson that the customer deals with, and more importantly, the zip code contribute in explaining the remaining half of the systematic variation in markups. In this subsection, I study how the seller responds to different dimensions of customer heterogeneity.

I first explore how sellers are responding to customer heterogeneity in the demand function. In the context of a wholesale market, as customers have different needs and are heterogeneous in their willingness to pay for each level of the intermediate good, one would expect quantity discounts to be part of the pricing decisions, as the sellers' tradeoff between markup and volume play out differently for different types of customers. Table 7 presents regressions of the markup on the percentile rank of the quantity purchased in units (i.e., a measure of how big each transaction is relative to the other transactions of the same good).[7] The table shows that if the percentile rank was an exogenous variable, there would be a negative effect of quantity on markups. The result is robust to controlling for customer and product fixed effects, and provides evidence that sellers respond to customer heterogeneity in demand.

Second, while one could think that zip code fixed effects affecting markups (see Table 6) is exclusively driven by the physical distance to the distributor's warehouse, Table 8 shows that although the physical distance is positively correlated with the markup level (but not statistically significant), it only explains a small share of the zip code effect. The remaining effect that is not captured by distance is likely to be capturing both demand shifters and rotators, as well as potential zip code-specific competition conditions.

In terms of how the seller is responding to heterogeneity in the extensive margin, or in how restaurants decide to break business relationships, I explore how markups are related to two proxies for search costs.

In Table 9 I study the correlation between markups and a first proxy for search cost. This proxy is defined as the fraction of the customer's total expenditure over time spent on the product being transacted.[8] One would expect that the higher the relative importance of a product in a customer's budget, the higher the benefits for the customer of obtaining a lower price, implying that the customer should be more price sensitive in the decision of moving her business elsewhere.[9] As the table shows, there is a negative correlation between the proxy for search cost and markup, which is consistent with sellers exploiting the higher market power that is given by a lower risk of losing a high-search cost customer. This relationship is robust to controlling for quantity, and customer and product fixed effects.

The second proxy for search costs is related to the fact that customers make multi-product purchases. One could think that it may be more costly to gather price quotes for a customer that buys all his inputs from a single food distributor than for one that purchases from several. The reason is that there may be economies of scale in gathering price quotes from a seller once the first price quote is requested. To explore this idea, I run regressions of markups on the number of other products that were purchased in the same purchase order (i.e., it is equal to zero if the customer bought only one product) as a fraction of the total number of products that were ever purchased by

[7] The benefit of using this measure is that it is comparable across products.

[8] That is, Customer's total expenditure on the product/Customer's total expenditure over time.

[9] An alternative way to interpret this variable is as a proxy for the propensity to search.

the customer (i.e., capturing the needs of a customer). Table 10 shows a robust positive correlation between markups and this proxy for search costs, which if exogenous, would imply that a customer that purchased everything from a single seller, would pay an average markup that is higher than one that does not. This is consistent with customers with lower search costs paying less, as they are more willing to search for a cheaper seller.

Summarizing, the first set of findings suggest that the seller responds to heterogeneity in demand. In particular, I find evidence that sellers tailor prices as a response to heterogeneity in how the volume-markup tradeoff plays out for different customers. The second set of findings suggest that the seller responds to customer behavior in regards to breakups, and that markups are positively correlated with two different proxies for search costs. Both sets of results suggest that the systematic markup variation is sensible to various sources of customer heterogeneity.

2.3.2 Price adjustments

Figure 4 shows the cumulative distribution functions of price changes, conditional on three different events. All three figures show, rather eloquently, that a bulk of the transactions were completed with no price change. That is, customers many times paid the same prices as in their previous transactions. This holds true even when conditioning on the event of a cost change (top left), and when conditioning on the event of a cost change of at least 5 percent (bottom).

Given the findings regarding price inertia in Figure 4, I explore when price changes occur. To study this, I consider the relationship between the decision of adjusting the price on two measures. The first is the cost change magnitude (as a percent of the the cost in the previous transaction). The second measure is the markup deviation, which I define by the difference between the markup that would result absent a price adjustment (i.e., using the price in the last transaction, and the current cost) and the average markup over time that was charged to the corresponding customer–product combination.[10] This second measure would be relevant if it was optimal to target a given markup to each customer.

The top row of Figure 5 shows that the probability of observing a price change increases with the magnitude of the cost change, independent of the direction of the cost change. Consistent with the price inertia discussion above, even for high cost changes, the probability of change is considerably less than one. The right-hand side graph shows that the cumulative distribution functions of the cost changes, restricted by type of price change, are ordered as expected according to the first-order stochastic dominance criterion. That is, the probability of observing a positive cost change together with a positive price change is greater than the probability of observing it with a null or negative price change.

[10] The first measure is defined as $(c_{jt} - c_{jt-1}/c_{jt-1})$, while the second as $(p_{ijt-1} - c_{jt})/p_{ijt-1}$ − mean markup$_{ij}$, where the subindeces i, j, t stand for customer, product, and time, respectively.

The bottom row of Figure 5 shows the same graphs but for the markup deviation measure. Again, the higher the deviation of the markup relative to the average customer-product markup, the higher the probability of observing a price change. As in the top row, the right-hand side graph shows that the cumulative distribution functions of the markup deviations, restricted by type of price change, are ordered as expected according to the first-order stochastic dominance criterion. That is, the probability of observing a positive markup deviation together with a negative price change is greater than the probability of observing it with a null or positive price change.

While there is evidence that the price change decision responds to markups that would be too far away from the average markup, and that prices respond to big cost changes, Figure 6 shows a smoothed plot of the expected price change conditional on both cost change and markup deviation. As expected, the graphs show that higher deviations from the average markup are observed together with lower price changes, and that higher cost changes are observed together with higher price changes. Both graphs also make clear the presence of price inertia, as the slope of the expected price change is relatively flat for both small cost changes and small markup deviations as many of these transactions are completed with null price changes.

These figures show that the price change decisions follow reasonable, and systematic patterns, and that the data is consistent with the idea that a price adjustment is made when the benefits exceed the costs of a price change. The higher the cost change or the higher the markup deviation, the higher the benefits of making a price adjustment. A price adjustment cost explains why sellers do not always adjust prices.

Table 11 shows the result of an exercise studying whether the salespeople are forward looking when adjusting prices. The table shows estimates of regressions where the dependent variable is a dummy for a price adjustment (positive and negative price change, respectively) conditional on the event of at least a 5 percent cost change (positive and negative cost change, respectively). The independent variable is a dummy that takes the value of one if in the following transaction there was also a cost change of at least 5 percent in the same direction. If it is costly to adjust a price, the price adjustment decision should be different when the price adjustment is motivated by a transitory cost increase as opposed to an upward trend in the wholesale cost. If salespeople are forward looking, the coefficient on the dummies for cost changes should be positive and significant. As the table shows, there is evidence in favor of salespeople exhibiting forward looking behavior for price increments.

Summarizing, I have shown that there is a non-trivial level of price inertia in these buyer–seller interactions, that the patterns of price adjustments are consistent with the existence of a price adjustment cost, and that there is evidence that salespeople internalize information that they may have regarding the evolution of cost in the future when making price adjustment decisions.

3 Model

In this section I formulate an industry-wide model of repeated purchases between sellers (food distributors) and customers (restaurants). The model studies the case of a single homogeneous good, as I find evidence of cross-product externalities in breakup decisions not playing a significant role (see Section 2.2). In the model, customers face search costs and make repeated purchases. Customers facing search costs gives sellers market power, as sellers can price above cost without losing a customer with certainty. Customers making repeated purchases affects the sellers' problem as sellers value preserving the future business of their customers. I assume sellers have learned the systematic differences across customers and set buyer-specific prices by trading off static profits with an increased likelihood of losing the customer's business (today and in the future). Price dispersion is a consequence of sellers practicing price discrimination, buyers facing search costs, and by cost heterogeneity across sellers.

3.1 Primitives of the model and timing of game

Consider an environment with N customers, and M sellers, where M is an arbitrarily large number. The sellers are ex-ante identical, and customers are heterogeneous. Each customer is characterized by a type $x \in X$. The sellers offer a homogeneous good that customers wish to purchase. Time is assumed to be discrete, and the horizon infinite.

Each period, customer i receives a price quote, p_{it}, from the seller with which she last did business with.[11] The customer decides whether to complete the transaction at the quoted price, or to instead break the relationship and search sequentially for a better price quote elsewhere. In terms of the timing of breakups (see the discussion in Section 2.2), I assume that all customers that decide to break a relationship, leave immediately without making further purchases from the seller.[12] The customer makes the breakup decision using i) her search cost, which determines how costly it is to obtain price quotes from alternative sellers; and, ii) her beliefs about the price distribution in the economy.

The fact that purchases are repeated affects the sellers' problem by adding dynamics to the pricing decision. Inertia in buyer–seller matches is a consequence of search costs and the assumption that customers get a free quote from the seller they transacted with last. As a consequence of rematching after a completed transaction, sellers take into account both the static and continuation value of serving a customer when deciding prices. While repeated transactions could bring along learning on the side of sellers, I assume that sellers have already learned the systematic differences across customer types.[13]

[11] The seller subindex was dropped for notational convenience.

[12] Assuming instead that they make a last purchase before leaving does not imply important changes to the comparative static results of the model.

[13] A richer model can include learning, but since the focus of the paper is understanding the implications of search

As a consequence of observing customer types, sellers set prices at the customer level, trading off the losses associated to a higher risk of losing a customer's business with the gains from serving the restaurant at a higher price. However, sellers face a price adjustment cost, that generates inertia with respect to the default price or the price of last transaction. Sellers set prices once they observe their wholesale cost, the customer-specific default price (i.e., the price of last period, $r_{it} \equiv p_{it-1}$), the price adjustment cost, and also given their beliefs about the customers' breakup strategies. I assume that each seller's wholesale cost is the sum of two components, a firm-specific shock, c_t, and a customer-seller specific shock, ν_{it}. The first component is meant to capture the cost of the good for the firm at time t, while the second, any additional cost that may affect the profitability of serving customer i in period t (e.g. traffic or delivery issues). c_t is assumed to be independent across sellers, while ν_{it} is assumed to be i.i.d. across sellers, customers, and time.

The relevant state variables each period are the firm-specific wholesale cost, c_t; the customer specific default price, r_{it}; the customer specific search cost level, ε_{it}; the customer-seller specific shock to the wholesale cost, ν_t; and the customer specific type, x_i.

The timing of the game is as follows for a customer-seller pair. At the beginning of each period, Nature makes two moves. First, it assigns the customer a search cost, $\varepsilon \in E$. Second, it draws the seller's period wholesale cost, $c \in C$, and the seller-customer specific shock to the wholesale cost, $\nu \in V$. Once nature has completed these moves, the seller proposes a price, $p \in P$, to the customer. This proposal is made after the seller observes the state variables (x, c, r, ν), and given his beliefs about the distribution of prices available in the economy, his beliefs about ε, and the price adjustment cost. The customer, having observed the proposal p, and the state variables (x, ε), decides whether to accept the price and complete the transaction, or to search for a better quote. The search or breakup decision is made as described above. If the customer decides to search for a better quote, then the customer gets randomly assigned to a new seller. The benefit of searching is that the customer may be matched with a lower cost seller, and thus, face a lower price (note that both components of the wholesale cost are assumed to be independent across sellers). Once rematched, the timing and stages of the game are just as above. The period is over when all customers have completed their purchase, either with the seller on hand, or with a new match.[14] Figure 7 summarizes the timing of the game, and gives a preview of the payoffs associated to each decision.

In terms of the distribution functions that govern the processes of the random variables discussed above, and the other relevant primitives of the model, I assume that

Assumption 1 (Primitives)

costs and price discrimination on welfare, adding learning implies another layer of complexity with little returns to the main purpose of the paper.

[14]For simplicity, I assume that when matched with a new seller, the seller is drawn a default price from the equilibrium price distribution with which he makes the pricing decision.

1. *The seller specific component of the wholesale cost, c, follows a Markov process, with CDF*

$$L : C^2 \to [0, 1],$$

where c is independent across sellers.

2. *The seller-customer specific component of the wholesale cost, ν, is an i.i.d. draw over time, and customer-seller combinations, from a distribution with CDF*

$$H : V \to [0, 1].$$

3. *The demand function of customers is given by,*

$$q : P \times X \to \mathbb{R}_{++},$$

which depends on price and on the customer type.

4. *The customer's search cost, ε, is an i.i.d. draw over time and customers, from a type–dependent distribution with CDF*

$$G : E \times X \to [0, 1].$$ [15]

The next subsections analyze, using backwards induction, the strategies that each customer and seller follow in equilibrium. The notion and existence of such equilibrium is also discussed below.

3.2 Customer's problem

Each period each customer starts her decision process after observing her type, x; her search cost, ε; and the price proposed by the seller, p. The customer also forms beliefs about the price distribution that she would encounter if she searches for price quotes with other sellers. Let these beliefs be type-specific and given by a CDF,

$$F : P \times X \to [0, 1].$$ [16]

Given the information that the customer observes, the payoff that the she obtains by accepting price p and completing the transaction is defined by,

$$\text{CS}(p; x) \equiv \int_p^\infty q(s; x) ds,$$

[15] While in Section 2.2 I argue that some patterns of the restaurants' behavior are consistent with a model where search costs are correlated over time, I assume for the purposes of the model that there is no such correlation. While allowing for correlation could potentially fit the data in a better way, it brings along modelling difficulties such as incorporating learning on the side of the seller which adds tradeoffs that are not central to the questions of this paper.

[16] In equilibrium, these beliefs reflect the pricing decisions of all sellers, and conditional on the customer type, the price dispersion comes from heterogeneity in cost across food distributors.

which is the customer surplus at price p.[17]

The customer, however, can also reject the price and search for a better quote. I assume that the search decision is sequential, and that customers cannot recall price quotes that were declined.[18] Conditional on searching, the customer is randomly assigned to a new seller. In this way, the expected payoff that the customer obtains by obtaining one additional price quote is given by,

$$\underbrace{\int \mathrm{CS}(s;x)dF(s;x)}_{\mathbb{E}[\text{Search reward}]} - \varepsilon,$$

which is the expected value of the customer's surplus (or the expected search reward) minus the search cost.

One can show that the optimal policy function that determines whether or not to search (i.e., break the relationship) is such that the customer rejects any price that is higher than a reservation price, $R(\varepsilon;x)$, and accepts anything else. This reserve price depends both on the customer's search cost and beliefs regarding the distribution of prices that are available in the economy. Alternatively, one can show that a myopic rule is also optimal. In this rule, the customer only needs to compare the net gains of obtaining one additional price quote with the payoff of accepting the price quote on hand. (DeGroot, 1970) That is, a customer facing state variables $(p,\varepsilon;x)$ will decide to search if and only if

$$\int \mathrm{CS}(s;x)dF(s;x) - \varepsilon \geq \mathrm{CS}(p;x). \tag{1}$$

Using (1), and given the distribution of the shocks to the search cost, the customer searches with probability,

$$\tilde{G}(p;x) = G\left(\int \mathrm{CS}(s;x)dF(s;x) - \mathrm{CS}(p;x);x\right), \tag{2}$$

where $G(\cdot;x)$ is the type–dependent CDF of ε given x. It is simple to check that $\tilde{G}(p;x)$ is increasing in p.

Finally, note that I assume that customers are myopic as they do not incorporate the existence of price inertia into their seller selection problem. Price inertia implies that being matched with a seller that quotes a high price today is not only unattractive in the current period, but also unattractive in the future as the same high price may be quoted again for the next transaction.

3.3 Seller's problem

Each seller offers a price, $p \in P$, to each customer he serves. Every offer is made after the seller observes the customer's type, x, the customer-specific default price, r, Nature's moves (c,ν), and

[17] Using Hotelling's lemma (see MWG), that relates the profit function to the supply correspondence at every price vector p, one can interpret the customer surplus at input price r as the cost savings associated to purchasing at price r relative to not purchasing at all (holding the others inputs fixed).

[18] Given that the policy function is of the reserve price type, this assumption is without loss of generality.

his beliefs about how the customer will react to every given price proposal. The seller also forms beliefs about the price distribution that is available to the customer if the customer breaks the relationship and searches for a better price elsewhere. Let these beliefs be given by a CDF,

$$F : P^2 \times C \times V \times X \to [0, 1],$$

which captures the joint distribution of prices, p, default prices, r, wholesale cost shocks, c and ν, as well as customer types, x.

The seller then solves

$$V(r, c, \nu; x) = \max \left\{ \max_{p \in P} \Lambda(p, c, \nu; x) - \eta, \ \Lambda(r, c, \nu; x) \right\}, \tag{3}$$

where $\Lambda(p, c, \nu; x)$ is the value of serving the customer at price p and state variables (c, ν, x), and where η is the price adjustment cost. The value function, V, captures the tradeoff between reoptimizing the price, and serving the customer with the default price but saving on the adjustment cost.

The value function of serving the customer at price p is given by,

$$\Lambda(p, c, \nu; x) = (1 - \tilde{G}(p; x)) \left(\pi(p, c, \nu; x) + \beta \mathbb{E}[V(p, c', \nu'; x)|c] \right) \tag{4}$$

where the period payoff function is given by

$$\pi(p, c, \nu; x) = q(p; x)(p - c - \nu),$$

and $(1 - \tilde{G}(\cdot; x))$ is the probability that a customer of type x will not search as a function of p (see (2)). I denote the solution to the seller's problem given in (3) as

$$p(r, c, \nu; x). \tag{5}$$

Conditional on deciding to adjust the price, the seller balances the tradeoff between increasing the price and achieving higher profits, with the losses associated to a higher risk of losing the customer's business today and in the future. This tradeoff depends on two margins. The first is the intensive margin, which is given by how much the customer will buy for every price conditional on accepting to do business with the seller. The second is the extensive margin, which is related to the price sensitivity of the customer's probability of breakup. The relevance of this second margin is given by the fact that a customer with a search cost distribution concentrated in high values will be less sensitive to prices when making her breakup decision. This low breakup elasticity becomes a source of market power, and gives the seller incentives to charge high markups as losing the customer's business (today and in the future) becomes less likely.

Other things to note in (4) are that the seller's continuation payoff depends on the price that the seller chooses in the current period, as it affects the price adjustment decision of the seller in the next period. As sellers are forward looking, they take into account the implications of the adjustment decision in the payoff of future periods as well.

Finally, the solution $p(r, c, \nu; x)$ is a function of both components of the wholesale cost. The unobserved shock to the wholesale cost, ν, generates price dispersion among customers facing the same observed state variables (i.e., r, c and x). When only conditioning on x, both components of the wholesale cost generate price dispersion among customers of the same type, x. Since both c and ν are i.i.d. across sellers, the type-specific price distribution is stationary. In equilibrium, the customers form their beliefs using the solution $p(r, c, \nu; x)$, and the distribution of its arguments.

3.4 Equilibrium

In this paper I use the notion of Markov perfect equilibrium, which in this context is defined as,

Definition 1 *An MPE is defined as a set of strategy functions (p^*, s^*), for sellers and customers, respectively, such that i) each seller is maximizing his value given the strategies of the other sellers and customers, ii) customers are making their search decisions, s^*, according to the rule given in (1), and iii) the beliefs of customers and sellers of how other sellers and customers behave are consistent with (p^*, s^*).*

The proof that at least one MPE exists is simple in the case of a discrete state space and set of actions. The proof uses the probability space representation of the MPE discussed in Aguirregabiria and Mira (2007). Define the equilibrium mapping, $P^* = \Lambda(\tilde{P})$, where each element of P_j^*, $j = 1, \ldots, M$, is given by

$$P_j^*(r, c; x) \equiv \int_{\nu \in V} 1\{p_j(r, c, \nu; x) = t\} dH(\nu),$$

where $p_j(r, c, \nu; x)$ is the solution to seller j's problem in (3), and where $\tilde{P}$ is the matrix that gathers the beliefs of each seller when solving his problem.

It is simple to check that the mapping $P^* = \Lambda(\tilde{P})$ is continuous in $\tilde{P}$, and since it maps a compact set of choice probabilities into itself, by Brouwer's fixed point theorem, there exists a fixed point to this mapping, guaranteeing the existence of at least one MPE.

3.5 Welfare and comparative statics

3.5.1 The welfare implications of price discrimination

To understand how price discrimination and search costs affect the pricing decision, I analyze a simple version of the model above in which there is no price adjustment cost, and the time horizon is of a single period. Sellers are heterogeneous in wholesale cost, c, and I assume search costs, ε, are distributed according to an exponential distribution with parameter λ. I assume that the demand function is given by $q(p) = \exp\{-\alpha p\}$. Customers are assumed to be heterogeneous in λ and α.

The seller's problem when allowed to price discriminate can be written as

$$\max_p O(p;\alpha,\lambda) \equiv \max_p \quad \exp\{-\lambda \max(\mathbb{E}[\text{Search reward}] - \text{CS}(p), 0)\} \exp\{-\alpha p\}(p-c), \qquad (6)$$

where the first term is the probability that the customer will not search if faced with price p (see definition in (2)), the second is the quantity demanded at price p, and the third is the price-cost margin. $\mathbb{E}$[Search reward] is the expected customer surplus that is obtained when searching exactly one seller, and is computed using the equilibrium price distribution.

The optimal pricing strategy with price discrimination is given by

$$p = \begin{cases} p^* \equiv c + 1/\alpha & \text{if } \text{CS}(p^*) > \mathbb{E}[\text{Reward}], \\ c + 1/(\lambda \exp\{-\alpha p^*\} + \alpha) & \text{otherwise}, \end{cases} \qquad (7)$$

which states that i) sellers set the monopoly price p^* if p^* is in the range of prices where not even the most efficient searchers have incentives to search and, ii) sellers balance the gains in profits with the losses associated to a decreased likelihood of serving the buyer if in the price range where there is search with positive probability (i.e., $\mathbb{E}[\text{Reward}] \geq \text{CS}(p^*)$).

Note that the pricing decision depends on the equilibrium search reward as the level of the equilibrium search reward determines which sellers are able to set the monopoly price. An interesting exercise is to analyze the effect of sellers deviating from a uniform price equilibrium by tailoring prices. Sellers tailoring prices will set prices according to rule (7), taking as given the expected search reward in the uniform price equilibrium, $\mathbb{E}[\text{Search reward}_U]$. Once all sellers adjust prices, buyers ajust their (type-specific) beliefs about the price distribution in the economy, and solve for the new equilibrium search reward, $\mathbb{E}[\text{Search reward}_{PD}]$. If $\mathbb{E}[\text{Search reward}_{PD,1}] < \mathbb{E}[\text{Search reward}_U]$, there is a set of sellers which were able to set the monopoly price p^* under $\mathbb{E}[\text{Search reward}_U]$, but are no longer able to do so under $\mathbb{E}[\text{Search reward}_{PD}]$.[19] The change in the expected search reward forces these sellers to revise their prices downwards as they are now in the price range where search occurs with positive probability and they are not balancing the tradeoff between profits and likelihood of serving the buyer. This downward adjustment of prices increases the competition for buyers and may create further price revisions. The process of revising prices because of the effect of competition continues until convergence of the type-specific price distribution (and, hence, until convergence of $\mathbb{E}[\text{Search reward}_{PD}]$).

Note also that the equilibrium search reward enters into the sellers objective function, implying that if the search reward has a sufficiently high increment, it can scale profits down to the point where the gains by price tailoring are offset by the increased competition across sellers.

The overall welfare implications of shifting from uniform pricing to price discrimination are given by netting i) the inefficient distribution of resources due to selling the good at different marginal

[19] A similar arguments applies to the case when $\mathbb{E}[\text{Search reward}_{PD,1}] > \mathbb{E}[\text{Search reward}_U]$.

valuations to different customers, ii) the overall effect on quantity, and iii) the effect of price discrimination on competition which affects both previous effects. The first two effects are standard in the discussion of price discrimination and welfare under the monopoly case (see Schmalansee (1981) or Varian (1985)), while the third, becomes relevant when studying a competitive environment (see Holmes (1989) or Corts (1998)). The difference of the framework presented here to other competitive frameworks that have been studied is that the effect of price discrimination on competition is given by how price discrimination changes search incentives.

Simulations of this version of the model suggest that price discrimination can be either welfare increasing or decreasing, and as discussed above, even decrease profits.

3.5.2 Comparative statics

Moving now the discussion to how customer heterogeneity affects the pricing decision of sellers, it is helpful to start analyzing the trade-off that sellers face and how this trade-off is affected by customer heterogeneity.

By taking the logarithm of the objective function, $O(p;\alpha,\lambda)$, and differentiating with respect to p, one obtains[20]

$$\frac{\partial \log O(p;\alpha,\lambda)}{\partial p} = -\alpha - \lambda \exp\{-\alpha p\} + \frac{1}{p-c}, \tag{8}$$

The derivative above captures the tradeoff that the seller faces when increasing the price. On the one hand, both the probability that the seller will serve the customer and the purchased quantity decrease (first two terms), and on the other, the seller obtains a higher margin (third term).

By differentiating (8) with respect to α and λ, one obtains

$$\frac{\partial^2 \log O(p;\alpha,\lambda)}{\partial p \partial \lambda} = -\exp\{-\alpha p\} < 0,$$
$$\frac{\partial^2 \log O(p;\alpha,\lambda)}{\partial p \partial \alpha} = -1 + \lambda p \exp\{-\alpha p\}.$$

The first equation shows that increasing λ (i.e., decreasing the mean of the distribution of search cost) implies that the customer's search decision becomes more sensitive to the price, which makes charging higher prices more costly for the seller. This effect implies that the sellers' optimal price is decreasing in λ.[21]

This comparative static result implies that the quantity distortions that customers face are a function of their elasticity in the search intensity margin. When $\lambda \to \infty$ (or when the search costs vanish), the price converges to the wholesale cost. When $\lambda \to 0$ (or when the search costs become arbitrarily large), the optimal price is the monopoly price.

[20] This condition requires that $\mathbb{E}[\text{Search reward}] - \text{CS}(p) > 0$, that is, it restricts to the range of prices for which there are values of ε for which it is profitable to search (or said differently, it restricts to the range of prices for which there is search with positive probability).

[21] This follows from Topkis' theorem and the fact that $O(p,\lambda)$ has increasing differences in $(p,-\lambda)$.

The effect of α on the sellers' tradeoff is not as straightforward as for λ because it has two conflicting effects. First, it intensifies the effect of the price increase on the quantity purchased, by making the customer purchase even less when faced with a price increase. Second, it increases the probability that the seller will serve the customer for any price p, as it reduces the customer sensitivity in this margin (note that $\lambda p \exp\{-\alpha p\}$ is the price elasticity of the probability of serving the customer). When

$$-1 + \lambda p \exp\{-\alpha p\} < 0,$$

that is, when the customer is relatively inelastic in the extensive margin, the first effect dominates, implying that p is decreasing in α (and vice versa).[22]

This analysis yields two testable implications: i) customers with lower search costs (or higher λ) face higher prices; and ii) when restricting attention to customers with high enough search costs (or low enough λ), the ones that are more price sensitive (higher α) face lower prices.

Note that these results are consistent with the descriptive exercises discussed in Subsection 2.3.1. For instance, Table 9 presents evidence that customers that have a higher proxy for search costs, are faced with higher markups.

3.5.3 Simulating the model

In this section I present numerical comparative static results of the full version of the model, that are meant to complement the discussion above, and shed light on identification. I present these results using a parameterization similar to the one that I use for the estimation.[23]

I assume that i) $P = [0 : 0.01 : 1.2]$, $C = \{0.15, 0.25, 0.35, 0.45, 0.55, 0.65, 0.75\}$; ii) the cost process exhibits persistence $\Pr(c_t = x | c_{t-1} = x) = 0.8$, and $\Pr(c_t = z | c_{t-1} = x) = 1/30$ for $z \neq x \in C$; iii) the unobserved shocks to the wholesale cost are such that $\nu \in [-.15, .15]$, with ν distributed according to a truncated normal distribution, $N(0, 2)$; iv) the demand function is given by, $q(p) = \exp\{1 - \alpha p\}$, with $\alpha =$ Price Coefficient; v) the search costs, ε, are distributed according to an exponential distribution, Exponential(λ), with parameter, λ, that is inversely related to the mean of the distribution, and with support $[-0.3, +\infty)$; and, vi) the price adjustment cost is given by $\eta \geq 0$.

Figure 8 show the results of the simulations. Consistent with the discussion above, Figure 8a shows that the effect of the price coefficient-demand is ambiguous, depending on the price sensitivity in the search decision margin. When $\lambda = 0$ (extreme case of price insensitivity in the search margin), customers with more price insensitive demand functions are faced with higher average markups. The opposite is true for higher values of λ as the effect of the price coefficient-demand on the search elasticity dominates that on quantity demanded (see discussion above).

[22]The formal argument is analogous to the one provided in Footnote 21.

[23]I would like to stress that I do not claim the generality of these results.

Figure 8b shows that higher values of the search cost parameter λ (or lower average search costs), imply lower average markups. That is, customers that are more sensitive to prices in the search or extensive margin, are favored with lower prices as they are more willing to break the relationship and move their business elsewhere. Interestingly, as displayed in Figure 8c, the price adjustment cost has no effect on the average markup.

Figure 8d shows that the price adjustment cost has a negative effect on the frequency of price adjustments, as expected. That is, the more costly it is to adjust the price, the fewer the price adjustments. Figure 8e shows that the search cost parameter, λ, intensifies the effect of the price adjustment cost on the frequency of price adjustments, but that is not able to generate infrequent price changes on its own.

These results clarify how the model will interpret the data. For a given demand function, the only parameter that affects markups is the search cost parameter. This implies that the variation in markups across types and over time is going to identify the search cost distribution as the model interprets this variation as heterogeneity in the search cost dimension. The frequency of price adjustments is affected by both the search cost parameter and the price adjustment cost, but is only possible for positive values of the latter. The heterogeneity in the frequency of price changes then identifies the price adjustment cost, taking into consideration how the (separately identified) search cost parameters affect price inertia.

4 Estimation, model specification, and results

4.1 Estimation

The estimation strategy that I use is a two-step procedure. The first step consists of estimating a subset of the primitives of the model. In the second step, I estimate the remaining primitives, using the first-step estimates, and by constructing a pseudo-likelihood function based on the equilibrium behavior of a seller.

For the estimation, I rely on the following assumption regarding the equilibrium that generates the data:

Assumption 2 (Estimation) *The data was generated by a stationary symmetric equilibrium.*

Under this assumption, it is sufficient to observe a single firm following the equilibrium strategy to identify equilibrium behavior and the primitives underlying the equilibrium behavior. The benefit of this assumption is that performing industry-wide inference is possible even when observing the behavior of a single firm.

The primitives I estimate in the first step are the wholesale cost process and the demand function. The remaining primitives to be estimated are i) the distribution of the search costs, $G : E \to \mathbb{R}$; ii) the price adjustment cost, η; and iii) the distribution of the customer-seller specific shocks to the wholesale cost, $H : V \to [0,1]$. I assume that G, η, and H are known functions up to a finite dimensional parameter $\omega_{01} \in \mathbb{R}^K$.

In the second step, I construct a pseudo likelihood function using the sellers' equilibrium behavior. I use the fact that the actions of a seller's rivals only enter into the seller's objective function through the expected search reward in the search probability,

$$\Pr(\text{Search}|p;x) = G(\mathbb{E}[\text{Search reward}|x] - \text{CS}(p);x).$$

As a consequence, it is sufficient to know the equilibrium search reward to compute the buyers' equilibrium search behavior.[24] Moreover, the best response of a seller to the equilibrium search behavior is also a best response to the equilibrium behavior of all other sellers. Hence, it is also sufficient to know the equilibrium search reward to compute the equilibrium behavior of the seller. Given that the data was generated by a stationary equilibrium, I treat the equilibrium search reward as a vector of type–dependent scalars to be estimated, ω_{02}. This approach simplifies the estimation of the game as it saves the econometrician from computing the equilibrium of the game for each trial parameter, $\omega = \{\omega_1, \omega_2\}$.

Using the first step estimates and the trial vector of parameters, ω, I compute the seller's best response to the search behavior and behavior of other sellers. Given the equilibrium search reward in the trial vector of parameters, ω_2, the seller's best response problem is a single-agent dynamic problem with a solution given by the function

$$p(r, c, \nu; x, \omega),$$

where r is the reference price, c is the observed wholesale cost, ν is the unobserved shock to the wholesale cost, x is the restaurant type, and ω the trial vector of parameters.

Using the function $p(r, c, \nu; x, \omega)$, I compute the conditional probabilities of observing a price t by integrating over the unobserved customer-seller specific shocks to the wholesale cost (ν), that is,

$$\Pr(p^{\text{model}} = t|r, c, x; \omega) = \int_{\nu \in V} 1\{p(r, c, \nu; x, \omega) = t\} dH(\nu, \omega), \tag{9}$$

where both the solution to the seller's problem, and the CDF H depend on the trial vector of parameters, ω, and where I make use that the set P is assumed to be discrete. By integrating out the unobserved shock to the wholesale cost, the model gives a distribution of prices conditional on the observed state variables: wholesale cost (c), reference price (r), and customer type (x).

Using the probabilities defined in (9), and that the data is organized in I tuples of transaction price, reference price, wholesale cost, and customer type, $\{p_i, r_i, c_i, x_i\}$, for each transaction $i =$

[24] Note that the demand function was consistently estimated in the first step.

$1, \ldots, I$, I define the pseudo log-likelihood function as

$$\mathcal{L}(\omega) = \sum_{i=1}^{I} \sum_{t \in P} 1\{p_i = t\} \log \left(\Pr(p^{\text{model}} = t; r_i, c_i, x_i, \omega) \right).^{25} \tag{10}$$

The pseudo-maximum likelihood estimate of ω_0 is defined as

$$\omega_{\text{PML}} = \underset{\omega \in \mathbb{R}^K}{\operatorname{argmax}} \, \mathcal{L}(\omega).$$

4.2 Model specification

The structural analysis is performed, separately, for two products. These products are highly transacted, and present differences in terms of markups, and frequency of price changes. The first product has a relatively low average markup, and a relatively high average frequency of price changes.[26] The second product, in change, has a relatively high average markup, and a low average frequency of price changes. While the fact that these items are highly transacted is econometrically convenient, the differences in their price outcomes is also beneficial as it will motivate a discussion in terms of understanding these differences in pricing outcomes from the point of view of the comparative statics of the model.

Regarding the specification of the primitives that are estimated in the first stage, I make the following assumptions,

Assumption 3 (Parametric assumptions I)

- $X = \{$*Pizza restaurant, Meat restaurant, Seafood restaurant, Other*$\}$, *that is, customer types will be captured by restaurant type.*
- *The demand function is given by* $q(p, \xi; x) = \exp\{\alpha_1 + \alpha_{2x} p + \xi\}$, *where the price slope is allowed to vary by customer type, and where* ξ *is an unobserved transaction-specific demand shock.*
- *The wholesale cost process,* $\Pr(c_t = c | c_{t-1} = c')$, *has the Markov property, and has support on a discrete set* C.

In terms of the estimation of these objects, I estimate the demand function using 2SLS. I use the wholesale cost, and both customer zip code and salesperson fixed effects as instruments. The wholesale cost process is estimated non-parametrically using sample frequency estimates.[27]

[25] I drop the first observation of each customer-product combination as I do not observe the default price, r, that the seller faced for these transactions.

[26] The relation is made with respect to the sample, not only with respect to the second product.

[27] That is, the estimate of transition probability from c to c' is given by

$$\hat{P}(c = c' | c_{-1} = c) = \frac{\sum_{j=1}^{N} 1\{c_j = c', c_{j,-1} = c\}}{\sum_{j=1}^{N} 1\{c_{-1} = c\}},$$

where c_j and $c_{j,-1}$ are the cost, and cost [lag] of observation j, respectively.

Regarding the specification of the primitives to be estimated by pseudo-maximum likelihood, I make the following assumptions,

Assumption 4 (Parametric assumptions II)

- *The price adjustment cost is assumed to be uniform across customers, and given by Adj. Cost =* η*.*

- *The search cost distribution is assumed to be an exponential distribution,*

$$\textit{Exponential}(\exp(-x'\mu)),$$

  with mean $\exp(x'\mu)$*.*

- *The equilibrium search reward is a type–dependent scalar,* $\gamma(x)$*.*

- *The unobserved shocks to the wholesale cost,* ν*, are assumed to be i.i.d. draws from a normal distribution* $N(0, \sigma_\nu^2)$*.*

Finally, time is re-scaled for each customer type according to their frequency of purchases, $T(x)$. That is, instead of using the discount factor β, I use $\beta^{T(x)}$. I set $\beta = 0.95$.

4.3 Identification

While I provided comparative statics results in Section 3.5 that help understand how the model would interpret the data, those results do not address how the model separately identifies the equilibrium expected search reward with the shape parameter of search cost distribution, $\lambda(x)$, and the price adjustment cost, η. In this subsection I will discuss identification when the equilibrium expected search is estimated as opposed to computed by imposing equilibrium conditions. For simplicity, I present analytic results for the static version of the model ($\beta = 0$) that extend to the dynamic version with slight modifications.

Define $R(0; x)$ as

$$\mathbb{E}[\text{Search reward}|x] = \text{CS}(R(0; x); x),$$

that is, as the reservation price of a customer of type x with a zero search cost. Note that no customers search for prices p such that $p < R(0; x)$, as not even the most efficient searcher (i.e., the customer with a zero search cost) finds searching attractive when quoted prices in this range.

Define the optimal price $p^*(c, \nu; x)$ as a function of the state variables $(c, \nu; x)$ as

$$p^*(c, \nu; x) = \underset{p}{\text{argmax}} \underbrace{\exp(-\lambda(x) \max\{0, \mathbb{E}[\text{Search reward}|x] - \text{CS}(p; x)\})}_{1-\Pr(\text{Search}|p;x)} \pi(p, c, \nu; x), \tag{11}$$

which when $p^*(c, \nu; x) > R(0; x)$ is given by

$$p^* = c + \nu + \frac{1}{\alpha_1(x) + \lambda(x) \exp\{\alpha_0 + \alpha_1(x) p^*\}}, \tag{12}$$

and when $p^*(c,\nu;x) < R(0;x)$,

$$p^* = c + \nu + \frac{1}{\alpha_1(x)}. \tag{13}$$

Note that the distinction of $p^* \gtreqless R(0;x)$ determines whether the search probability is positive or zero at p^*.

Define the value function of the seller when facing state variables $(c,r,\nu;x)$ as

$$\begin{aligned} V(r,c,\nu;x) = \max\Bigg\{ & \exp(-\lambda(x)\max\{0, \mathbb{E}[\text{Search reward}|x] - \text{CS}(r;x)\})\pi(r,c,\nu;x), \\ & -\eta + \max_p \{\exp(-\lambda(x)\max\{0, \mathbb{E}[\text{Search reward}|x] - \text{CS}(p;x)\})\pi(p,c,\nu;x)\}\Bigg\}, \end{aligned} \tag{14}$$

which takes into account the price adjustment decision of whether to adjust the price to p^* or to serve the customer at the default price r.

In what follows I study two pricing situations that illustrate how i) the adjustment cost, η, ii) the shape parameter of the search cost distribution, $\lambda(x)$, and iii) $\mathbb{E}[\text{Search reward}|x]$ affect the sellers' problem differently depending on the pricing situation. The fact that the parameters affect the sellers' problem differently in these different pricing situations allows to separately identify these parameters.

Case 1: $p^*(c,\nu;x) < R(0;x)$ and $r < R(0;x)$.

In this case, the seller faces no risk of losing the customer. The value function in (14) can be expressed as,

$$V(r,c,\nu;x) = \max\left\{\pi(r,c,\nu;x), -\eta + \pi(p^*(c,\nu;x),c,\nu;x)\right\}. \tag{15}$$

Note that $\lambda(x)$ (parameter of the search cost distribution) and $\mathbb{E}[\text{Search reward}|x]$ do not affect the pricing decision of the seller in this case because the relevant prices are in the price range where there is no search (see (13)). The price adjustment cost is identified through the extent up to which the seller sets $p = r \neq p^*(c,\nu;x)$.

Case 2: $p^*(c,\nu;x) > R(0;x)$.

In this case, the seller faces the risk of losing the customer when making a price adjustment. The value function in (14) can be expressed as,

$$\begin{aligned} V(r,c,\nu;x) = \max\Bigg\{ & \exp(-\lambda(x)\max\{0, \mathbb{E}[\text{Search reward}|x] - \text{CS}(r;x)\})\pi(r,c,\nu;x), \\ & -\eta + \exp(-\lambda(x)(\mathbb{E}[\text{Search reward}|x] - \text{CS}(p^*(c,\nu;x);x)\}))\pi(p^*(c,\nu;x),c,\nu;x)\Bigg\}. \end{aligned}$$

Note that $\lambda(x)$ (parameter of the search cost distribution) affects the adjusted price p^*, while $\mathbb{E}[\text{Search reward}|x]$ does not (see (12)). $\mathbb{E}[\text{Search reward}|x]$, however, affects the decision of whether to adjust the price, as it shifts the level of the probability of serving the customer at every price $p \geq R(0;x)$. $\lambda(x)$ is separately identified from $\mathbb{E}[\text{Search reward}|x]$ as the former affects the optimal price conditional on making a price adjustment, while the latter only affects the decision of whether to adjust the price. Given that the price adjustment cost, η, is identified from the pricing situations in Case 1, $\mathbb{E}[\text{Search reward}|x]$ and $\lambda(x)$ are identified by Case 2–variation in both the price adjustment decisions and the optimal prices conditional on making a price adjustment.

Finally, it is important to note that the estimate of $\mathbb{E}[\text{Search reward}|x]$ would be affected by search costs having a a non-zero location parameter. Using the estimated model, one can compute the equilibrium $\mathbb{E}[\text{Search reward}|x]$ and contrast it with estimated value. Any difference between the prediction of the model and the estimate is given by a non-zero location parameter,

$$\text{Location}(x) = \mathbb{E}[\text{Search reward}|x]^{\text{Model}} - \mathbb{E}[\text{Search reward}|x]^{\text{Estimate}}.$$

Note that the existence of a non-zero location parameter does not affect the identification of the shape parameter of search costs nor the price adjustment cost.

4.4 Estimates

Table 12 shows the first-step estimates of the demand function. The 2SLS estimates show a strong negative effect of price on quantity, and show heterogeneity at the customer level. Table 14 shows that the implied price elasticities of demand (computed at the customer type average price) take reasonable magnitudes, and are uniformly smaller for the high-markup product compared to the low-markup product. The estimates imply that a one percent increase in price leads to an average 2.94 percent decrease in quantity for the low-markup product, and to a 1.82 percent decrease in quantity for the high-markup product.

Table 13 shows the results of the pseudo maximum likelihood estimation. The search cost parameter estimates show heterogeneity at the customer level, and as can be seen in Table 15, they imply large price elasticities in the extensive margin. On average, the customers are more elastic in the low-markup product relative to the high-markup product. The estimates imply that a one percent increase in price leads to an average 10.35 percent decrease in the probability of serving the customer for the low-markup product, and to a 8.18 percent decrease for the high-markup product.

Figure 9 show the equilibrium reservation prices as a function of the search cost (see Section 3.2 for a formal discussion). These graphs capture several aspects of the behavior of customers. First, the intercept in the vertical axis when the search cost is zero gives the upper bound of the region of prices for which customers never search. For this region of prices, not even the customers that are the most efficient searchers (i.e., null search cost) find it worthwhile to search. The price elasticity of the search probability is zero in this range of prices, as no price in this region induces search. The upper bound of the no-search price region exhibits heterogeneity across customer types. In the high markup product, the Other-customer type never searches for prices below \$14.5, while the Pizza-customer type never searches for prices below \$10.5. Second, the graphs show heterogeneity in the level of reservation prices. The Seafood-type is the customer type with the lowest reservation prices for each level of search cost in the low-markup product, while for the high-markup product, it is the Pizza-type. Third, the price elasticity of the search probability in the range of prices for which there is search is given by the slope of each reservation price curve. The slope of the reservation price curves display heterogeneity that predict important differences in the behavior of customers. In the low-markup product, a seller loses the average Pizza-type customer (given by the horizontal position of the marker in the corresponding curve) by increasing the price from $25 to $27.5, while the seller loses the average Seafood-type customer by only increasing the price from $24.5 to $25.5.

The adjustment cost parameter (i.e., the logarithm of the adjustment cost) takes low values for both products (only cents of a dollar). To measure the implications of these estimates, I contrast the price adjustment frequency in the data with the predictions of the estimated models. The results are displayed in Table 16, and as can be seen, the model predicts an important amount of price inertia. For all except one case, the model over predicts the frequency of price adjustments, but correctly predicts that price inertia is on average higher in the high-markup product than in the low-markup one.

Finally, the unobserved heterogeneity in the wholesale cost is found to play a similar role for both products in terms of helping the model fit the data.

5 Search costs, welfare, and competition

In this section I use the estimates of the structural model for two purposes. First, I quantify market power and the overall welfare losses generated by search costs. Second, I answer the theoretically ambiguous question of how price discrimination affects welfare and competition in the context of the market under study.

The welfare measure I use is the expected total (customer) surplus of a completed buyer–seller interaction. This measure incorporates the selection process on the side of buyers, as the expectation is conditional on that the customer accepts to complete the transaction at the quoted price. I compute the welfare measure under four scenarios: i) the price discrimination equilibrium, ii) the uniform price equilibrium (i.e., sellers set uniform prices), iii) frictionless perfect competition (i.e., no welfare losses), iv) the monopoly case (i.e., limit case of current market structure when search costs are arbitrarily large).

5.1 Measuring the welfare implications of search costs

In this subsection, I use the structural estimates to measure market power and the welfare implications of search costs. To implement this analysis I compare the current market structure to efficient perfect competition and the monopoly case. Note that both perfect competition and the monopoly case are limit cases of the current market structure when search costs are arbitrarily large and small, respectively. Perfect competition and monopoly give upper and lower bounds, respectively, to total surplus in this market. Comparing the current market structure with these two benchmarks allows to quantify how much welfare is lost relative to perfect competition, and measure how these welfare losses compare to the welfare losses when the market shifts from perfect competition to being monopolized. In addition, the comparison allows to measure how much of the monopoly profits sellers are able to extract with the existence of search costs.

To compute the welfare measures of a buyer–seller interaction under the current market structure, I use the price discrimination equilibrium solution, $p(r, c, \nu; x)$, discussed in Section 3.3. As discussed above, the welfare measure I compute is the expected total surplus of a completed buyer–seller interaction, which requires the price to be below the type-specific reservation price,

$$p \leq R(\varepsilon; x).$$

In this way, the welfare measures (i.e., expected customer and total surplus, respectively) of the current market structure are given by

$$\begin{aligned} \mathbb{E}[\mathrm{CS}_{\mathrm{PD}}(x)] &= \mathbb{E}_{\varepsilon}\left[\mathbb{E}\left[\mathrm{CS}(p(r, c, \nu; x); x) | p \leq R(\varepsilon; x)\right]\right], \\ \mathbb{E}[\mathrm{Profits}_{\mathrm{PD}}(x)] &= \mathbb{E}_{\varepsilon}\left[\mathbb{E}\left[(p(r, c, \nu; x) - c - \nu) q(p(r, c, \nu; x); x) | p \leq R(\varepsilon; x)\right]\right], \end{aligned} \tag{16}$$

where the expectation is taken first over the state variables (r, c, ν) using the type–dependent price discrimination equilibrium joint distribution of these variables, and second, over search costs using the type–dependent estimated distribution of search costs.

The welfare measures under the perfect competition market structure give an upper bound to total surplus. Under this market structure, I assume that the lowest-cost firm serves all customers. That is, all customers always face the price $p^C = \underline{c} + \underline{\nu}$. The welfare measures under perfect competition are given by

$$\begin{aligned} \mathbb{E}[\mathrm{CS}(x)]_{\mathrm{PC}} &= \mathrm{CS}(\underline{c} + \underline{\nu}; x), \\ \mathbb{E}[\mathrm{Profits}(x)]_{\mathrm{PC}} &= 0. \end{aligned}$$

The welfare measures under the monopoly case give the lower bound to total surplus, and give a measure of how total surplus would change if all sellers were to merge into a single firm. Given that a monopolist faces no competition, the monopolist is unconstrained in setting the customer-specific monopoly price, $p^{\mathrm{M}}(c, \nu; x)$.[28] The welfare measures under monopoly are given by

$$\begin{aligned} \mathbb{E}[\mathrm{CS}(x)]_{\mathrm{M}} &= \mathbb{E}\left[\mathrm{CS}(p^{\mathrm{M}}(c, \nu; x); x)\right], \\ \mathbb{E}[\mathrm{Profits}(x)]_{\mathrm{M}} &= \mathbb{E}\left[(p^{\mathrm{M}}(c, \nu; x) - c - \nu) q(p^{\mathrm{M}}(c, \nu; x); x)\right], \end{aligned}$$

[28]Note that I assume that the monopolist adjusts prices without friction.

where the expectation is taken over the state variables (c, ν), using the estimated distributions of these variables.

Using the welfare measures, I perform three comparisons. First, to obtain a measure of how much welfare is lost with customers facing search costs, I compare total surplus under the current market structure to total surplus under perfect competition. Second, I compare the loss in welfare generated by search costs relative to perfect competition to the loss in welfare when shifting from perfect competition to monopoly. This second measure allows to put the search cost implications on welfare in perspective with the more severe welfare losses when going from perfect competition to monopoly. Third, to obtain a measure of market power and a measure of how competition affects welfare in presence of search costs, I compare profits and total surplus under the current market structure relative to monopoly.

Finally, I compute the level of marginal cost savings that would be needed for a monopolist to replicate the total surplus under the current market structure. This gives a measure of how much gains in efficiencies would be required to compensate the welfare loss due to the lack of competition under monopoly. For this exercise, I calculate the welfare measures in the case of a monopolist with marginal costs scaled down by a factor of τ,

$$\begin{aligned}
\mathbb{E}[\mathrm{CS}(x)|\tau]_{\mathrm{M}} &= \mathbb{E}\left[\mathrm{CS}(p^{\mathrm{M}}(\tau c, \nu; x); x)\right], \\
\mathbb{E}[\mathrm{Profits}(x)|\tau]_{\mathrm{M}} &= \mathbb{E}\left[(p^{\mathrm{M}}(\tau c, \nu; x) - \tau c - \nu) q(p^{\mathrm{M}}(\tau c, \nu; x); x)\right].
\end{aligned}$$

The minimum cost reduction required for the merger to be welfare enhancing is given by finding the lowest value of τ for which the inequality,

$$\mathbb{E}[\mathrm{CS}(x)|\tau] + \mathbb{E}[\mathrm{Profits}(x)|\tau]_{\mathrm{M}} \geq \mathbb{E}[\mathrm{CS}_{\mathrm{PD}}(x)] + \mathbb{E}[\mathrm{Profits}_{\mathrm{PD}}(x)],$$

is satisfied.

5.1.1 Results

Table 17 shows the results of comparing the customer surplus, profits, and total surplus of a completed buyer–seller transaction, under different market structures. The results are reported as a fraction of the total surplus under perfect competition.

To quantify the overall welfare implications of the market power generated by search costs, I compare the customer and total surplus under the current market structure to the total surplus under perfect competition (columns 2 and 5, respectively). The loss in welfare is substantial for both products, but lower for the low-markup product, where customers are on average more price elastic in the extensive margin or search probability (see Table 15). In the low-markup product, the overall loss in total surplus is almost 36 percent of the total surplus under perfect competition, and customers only capture 58 percent of the perfect-competition customer surplus. The loss in welfare is lowest (both, in customer and total surplus) for the Seafood-customer type, that has the lowest reservation price for every search cost (see Figure 9) and that is the most price-elastic in the extensive margin. In the high-markup product, customer and total surplus decrease by about 50 and 42 percent, respectively, relative to the total surplus under perfect competition. Just as in the low-markup product, the loss in welfare (in both measures) is lowest for the most elastic customer type in the extensive margin (i.e., Pizza).

To measure market power, I compare the profits with search costs to the profits under monopoly (columns 4 and 1, respectively). In the high-markup product, sellers capture 35.58 percent of the monopoly profits while in the low-markup product, only 33.73 percent. That rent extraction is higher in the high-markup product is consistent with buyers of the high-markup product being on average less price elastic both in demand and in search probability. The share of monopoly profits that sellers obtain is highest for the customer types that have the lowest extensive margin price elasticities for each product. In the case of the high-markup product, rent extraction reaches 66.56 percent of monopoly profits for the Seafood restaurant type, while in the low-markup product, it reaches 46.30 percent of the monopoly profits for the Pizza restaurant type.

To quantify the role of competition under the current market structure, I compare the welfare measures under the current market structure to total surplus under monopoly (column 2). The table shows that shifting from the current market structure to monopoly would reduce total surplus in almost 50 percent, implying that competition plays a significant role limiting rent extraction under the current market structure. The comparison with the monopoly case also gives a way of putting the losses generated by search costs in perspective. The welfare losses generated by search costs relative to perfect competition for the low and high markup product, respectively, are 60 and 65 percent of the losses from shifting from perfect competition to monopoly.

Finally, I simulate by how much a monopolist would have to reduce marginal costs to compensate the loss in total surplus generated by eliminating competition. Table 18 shows the results of this exercise. I find that a monopolist in the low-markup product that reduces marginal costs by 18 percent compensates the welfare loss in total surplus due to the lack of competition when monopolizing the market. In the high-markup product, this number raises to 30 percent. This analysis abstracts away from any potential savings in fixed costs of operation.

5.2 Measuring the welfare implications of price discrimination

In this second counterfactual, I analyze both the welfare and the competition implications of sellers practicing price discrimination, by comparing the equilibrium when sellers tailor prices to the equilibrium when sellers set uniform prices.

Similar to Section 3.3, the sellers' problem when facing customer type uncertainty is given by

$$V(r, c, \nu) = \max\{\Lambda(r, c, \nu), \ \ \max_p \Lambda(p, c, \nu) - \eta\},$$

where η is the price adjustment cost, and where Λ is the expected value of serving a customer at price p,

$$\Lambda(p, c, \nu) = \left\{ \sum_{z \in X} \mu(z) \left[(1 - \tilde{G}(p; z)) \left(\pi(p, c, \nu; z) + \beta \mathbb{E}[V(p, c', \nu')|c] \right) \right] \right\},$$

where $\mu(z)$ is the probability that the seller will encounter a customer of type z. By definition,

$$\sum_{x \in X} \mu(x) = 1.$$

I denote the uniform-price equilibrium pricing rule by

$$p^{\text{No PD}}(r, c, \nu).$$

The problem of the seller under the price discrimination case is the one described in Section (3.3). I denote the price discrimination equilibrium pricing rule by

$$p^{\text{PD}}(r, c, \nu; x).$$

To compute how welfare changes when moving from one equilibrium to the other, I compute the expected customer and seller surplus of a completed buyer–seller interaction under both equilibria. As discussed above, this welfare measure takes into account the selection process on the side of the buyer. For a buyer to accept to transact with a seller, the quoted price must be below the type-specific reservation price,

$$p \leq R(\varepsilon; x).$$

The welfare measures under equilibrium $k \in \{\text{Price discrimination}, \text{No price discrimination}\}$ are then given by

$$\begin{aligned} \mathbb{E}[\text{CS}_k(x)] &= \mathbb{E}_\varepsilon \left[\mathbb{E} \left[\text{CS}(p^k(r, c, \nu; x); x) | p \leq R(\varepsilon; x) \right] \right], \\ \mathbb{E}[\text{Profits}_k(x)] &= \mathbb{E}_\varepsilon \left[\mathbb{E} \left[(p^k(r, c, \nu; x) - c - \nu) q(p^k(r, c, \nu; x); x) | p \leq R(\varepsilon; x) \right] \right], \end{aligned} \tag{17}$$

where the expectation is taken first over the state variables (r, c, ν) using the type–dependent joint distribution of these variables under equilibrium k, and second, over search costs using the type–dependent estimated distribution of search costs, and where $p^k(r, c, \nu; x)$ is the policy function under equilibrium k.[29]

5.2.1 Results

Table 19 shows the results of analyzing how equilibrium profits and total surplus change with price discrimination, measured relative to the corresponding welfare measures under the uniform price

[29] Note that $p^{\text{No PD}}(r, c, \nu; x) = p^{\text{No PD}}(r, c, \nu)$, for all x.

equilibrium. The table also reports how the customers' equilibrium expected search reward changes when food distributors shift from setting uniform prices to practicing price discrimination.

The first column of each panel shows how the expected search reward changes when shifting from the uniform price equilibrium to the price discrimination equilibrium. On average, the expected search reward increases for both products, implying an increase in the propensity to search and, hence, in intensified competition. The increase in competition intensity is higher in the high-markup product. Although the average change in the search reward is no more than 2 percent, the search reward can change by almost 11 percent for a customer type-product combination. For both products, competition decreases the most for the customer types that are the least price elastic in the (seller substitution) extensive margin (see Tables 14 and 15). As these customers are less inclined to search, sellers raise the prices to these customers, lowering the incentives to search even more.

The second column shows how customer surplus changes when shifting from one equilibrium to the other. Customer surplus increases overall by no more than 2 percent. Whenever the search reward or competition intensity increases, customer surplus also increases as a consequence of customers facing a more favorable equilibrium price distribution. Note that customer surplus decreases only for the customer types that are the least price elastic in the extensive margin for each product, where as discussed above, competition is softened.

The effect of price discrimination on profits varies (see third column). Overall profits decrease for the low-markup product, but increases in more than 50 percent for the high-markup product. In the low-markup product, profits increase the most for the customer type that is the most price inelastic in the extensive margin and most price elastic in the intensive margin. Note that competition is softened for this customer type, implying that sellers not only benefit from price tailoring, but also from less competition. In the high-markup product, profits increase the most for the customer types that are the most (i.e., Pizza) and least price elastic (i.e., Seafood), respectively, in the extensive margin. It is interesting to note that profits increase for the Pizza-type despite intensified competition.

Total surplus increases overall for both markets (see last column), making price discrimination a welfare increasing practice. The positive effect is modest in the low-markup product, reaching less than a 1 percent increase in total surplus relative to the uniform-price equilibrium. In the low-markup product, total surplus increases despite a decrease in equilibrium profits. In the case of the high-markup product, on average both customer surplus and profits increase, which make total surplus increase by almost 6 percent.

Summarizing, these results highlight that the ability to tailor prices can have an important impact on profits if there is enough customer heterogeneity and if it does not come along with an increased competition. They also show that price discrimination is overall welfare increasing, which is partially due to price discrimination intensifying competition.

6 Discussion

Many markets for homogeneous goods present heterogeneity in the prices paid by buyers. One explanation of how price dispersion can be sustained in equilibrium is that buyers face search costs. Search costs are a concern because they reduce efficiency by limiting customer mobility to low-price sellers. However, search costs are not the only factor that can cause price heterogeneity and that affects efficiency. In markets where sellers choose not to commit to posted prices, price dispersion across sellers may coexist with price discrimination.

I study a wholesale market for commodity goods with an important extent of product-specific price and markup dispersion. On average, about two-thirds of the product-specific markup variation is between-customer variation, providing evidence that price discrimination is actively practiced. At the same time, buyers respond to prices along two margins. The quantity purchased is price sensitive, and the higher the price the more likely that a buyer will break a business relationship. I explain the data using a search model. Search costs give sellers market power as search costs reduce the propensity of buyers to search for low-price firms. Sellers set buyer-specific prices by trading off static profits with the change in the expected value of serving a customer.

The objective of this paper is twofold. First, I measure market power and the welfare implications of search costs. Second, I solve the empirical question of how price discrimination affects welfare and competition in this market. Price discrimination not only affects quantity and generates allocation inefficiencies, but may also intensify competition depending on how it changes search incentives.

I find that sellers are able to extract almost 36 percent of the monopoly profits using the search cost market power. The loss in total surplus generated by search costs is significant, and can be as high as 42 percent of the total surplus under perfect competition. The losses in total surplus generated by search costs represent two-thirds of the losses in welfare when shifting from perfect competition to monopoly. I find that while competition plays a relevant role in the current market structure with search costs, monopolizing the market can be welfare enhancing if the monopolist reduces marginal cost by 30 percent (leaving aside potential savings in fixed costs of operation).

I find that price discrimination increases total surplus relative to the case where sellers set uniform prices, and that this result is partially explained by price discrimination increasing search incentives, and hence, intensifying competition. While it is always a dominating strategy for sellers to price discriminate, the effect of price discrimination on intensifying competition is strong enough that it makes equilibrium profits decrease for some customer type-product combinations.

On the methodological side, this paper contributes a framework for analyzing competition and mergers in markets for homogeneous goods where buyers face search costs. In this context, the elasticity of substitution has to be reinterpreted as the elasticity of the search probability, which to be computed, requires estimating a search model. Standard discrete choice demand models fail to capture that the elasticity of the search probability is completely inelastic for prices where buyers never search, and hence, underestimate market power. This paper also contributes a new way of identifying search costs. The identification strategy relies on the sellers' decision of trading off static profits with the change in the expected value of serving a customer.

References

Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75(1), 1–53.

Aguirregabiria, V., and A. Nevo (2013): "Recent Developments in Empirical IO: Dynamic Demand and Dynamic Games," *Proceedings of the Econometric Society World Congress*, Forthcoming.

Anderson, E. T., and D. I. Simester (2010): "Price stickiness and customer antagonism," *Quarterly Journal of Economics*, 125(2), 729–765.

Armstrong, M., and J. Vickers (2001): "Competitive Price Discrimination," *The RAND Journal of Economics*, 32(4), 579–605.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Prominence and consumer search," *The RAND Journal of Economics*, 40(2), 209–233.

Borenstein, S. (1985): "Price Discrimination in Free-Entry Markets," *The RAND Journal of Economics*, 16(3), 380–397.

Borenstein, S., and N. L. Rose (1994): "Competition and Price Dispersion in the U.S. Airline Industry," *Journal of Political Economy*, 102(4), 653–683.

Borenstein, S., and A. Shepard (2002): "Sticky prices, inventories, and market power in wholesale gasoline markets," *The RAND Journal of Economics*, 33(1), 116–139.

Brynjolfsson, E., and M. Smith (2000): "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Marketing Science*, 46, 563–585.

Burdett, K., and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51(4), 955–969.

Busse, M., and M. Rysman (2005): "Competition and price discrimination in Yellow Pages advertising," *The RAND Journal of Economics*, 36(2), 378–390.

Chandra, A., and M. Tappata (2011): "Consumer search and dynamic price dispersion: an application to gasoline markets," *The RAND Journal of Economics*, 42(4), 681–704.

Chevalier, J. A., A. K. Kashyap, and P. E. Rossi (2003): "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *The American Economic Review*, 93(1), 15–37.

Clay, K., R. Krishnan, and E. Wolff (2001): "Prices and Price Dispersion on the Web: Evidence from the Online Book Industry," *Journal of Industrial Economics*, 49, 521–540.

Corts, K. S. (1998): "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," *The RAND Journal of Economics*, 29(2), 306–323.

Crawford, G. S., and A. Yurukoglu (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *The American Economic Review*, 102(2), 643–685.

Dafny, L. S. (2010): "Are Health Insurance Markets Competitive?," *The American Economic Review*, 100, 1399–1431.

De los Santos, B., A. Hortacsu, and M. R. Wildenbeest (2012): "Testing Models of Consumer Search Behavior Using Data on Web Browsing and Consumer Purchases," *The American Economic Review*, 102(6), 2955–2980.

DeGroot, M. H. (1970): *Optimal Statistical Decisions*. McGraw-Hill, Inc.

Diamond, P. A. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 156–168.

Eckstein, Z., and K. Wolpin (1990): "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, 58, 783–808.

Ellison, G., and S. F. Ellison (2009): "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77(2), 427–452.

Goldberg, P. K. (1996): "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 104(3), 622–654.

Goolsbee, A. (2001): "Competition in the Computer Industry: Online vs. Retail," *Journal of Industrial Economics*, 49, 487–499.

Gowrisankaran, G., A. Nevo, and R. Town (2013): "Mergers when Prices are Negotiated: Evidence from the Hospital Industry," Manuscript. Northwestern University.

Grennan, M. (2013): "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *The American Economic Review*, 103(1), 145–177.

Heidhues, P., and B. Koszegi (2008): "Competition and Price Variation when Consumers are Loss Averse," *The American Economic Review*, 98, 1245–1268.

HENDEL, I., AND A. NEVO (2011): "Intertemporal Price Discrimination in Storable Goods Markets," Manuscript. Department of Economics, Northwestern University.

HOLMES, T. J. (1989): "The Effects of Third-Degree Price Discrimination in Oligopoly," *The American Economic Review*, 79(1), 244–250.

HONG, H., AND M. SHUM (2006): "Using price distributions to estimate search costs," *The RAND Journal of Economics*, 37(2), 257–275.

HORTACSU, A., AND C. SYVERSON (2004): "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds," *Quarterly Journal of Economics*, 119(1), 403–456.

KASHYAP, A. K. (1995): "Sticky Prices: New Evidence From Retail Catalogs," *Quarterly Journal of Economics*, 110(1), 245–274.

KATZ, M. L. (1984): "Price Discrimination and Monopolistic Competition," *Econometrica*, 52(6), 1453–1471.

——— (1987): "The Welfare Effects of Third-Degree Price Discrimination in Intermediate Goods markets," *The American Economic Review*, 77(1), 154–167.

KENISTON, D. (2011): "Bargaining and Welfare: A Dynamic Structural Analysis of the Autorickshaw Market," Manuscript. Yale University.

KOSZEGI, B., AND M. RABIN (2006): "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121, 1133–1166.

LESLIE, P. (2004): "Price Discrimination in Broadway Theater," *The RAND Journal of Economics*, 35(3), 520–541.

LEWIS, M. S. (2010): "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market," Manuscript. The Ohio State University.

MARSHALL, G. (2013): "Hassle Costs and Price Discrimination: An Empirical Welfare Analysis," Manuscript. Northwestern University.

MCCALL, J. J. (1970): "Economics of Information and Job Search," *Quarterly Journal of Economics*, 84, 113–126.

MORTIMER, J. H. (2007): "Price Discrimination, Copyright Law, and Technological Innovation: Evidence from the Introduction of DVDs," *Quarterly Journal of Economics*, 122(3), 1307–1350.

Nakamura, E., and J. Steinsson (2011): "Price setting in forward-looking customer markets," *Journal of Monetary Economics*, 58(3), 220–233.

Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," *The Review of Economic Studies*, 75, 901–928.

Reinganum, J. F. (1979): "A Simple Model of Equilibrium Price Dispersion," *Journal of Political Economy*, 87(4), 851–858.

Renner, E., and J.-R. Tyran (2004): "Price rigidity in customer markets," *Journal of Economic Behavior & Organization*, 55(4), 575–593.

Rob, R. (1985): "Equilibrium Price Distributions," *The Review of Economic Studies*, LII, 487–504.

Rosenthal, R. W. (1982): "A Dynamic Model of Duopoly with Customer Loyalties," *Journal of Economic Theory*, 27, 69–76.

——— (1986): "Dynamic Duopoly with Incomplete Customer Loyalties," *International Economic Review*, 27(2), 399–406.

Salop, S. (1977): "The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination," *The Review of Economic Studies*, 44(3), 393–406.

Schmalansee, R. (1981): "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination," *The American Economic Review*, 71(1), 242–247.

Shepard, A. (1991): "Price Discrimination and Retail Configuration," *Journal of Political Economy*, 99(1), 30–53.

Sorensen, A. T. (2000): "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, 108(4), 833–850.

Spiegler, R. (2012): "Monopoly Pricing when Consumers are Antagonized by Unexpected Price Increases: A "Cover Version" of the Heidhues-Koszegi-Rabin Model," *Economic Theory*, 51, 695–711.

Stahl, D. O. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *The American Economic Review*, 79, 700–712.

Tappata, M. (2009): "Rockets and feathers: Understanding asymmetric pricing," *The RAND Journal of Economics*, 40(4), 673–687.

Varian, H. R. (1980): "A Model of Sales," *The American Economic Review*, 70(4), 651–659.

——— (1985): “Price Discrimination and Social Welfare,” *The American Economic Review*, 75(4), 870–875.

Villas-Boas, S. B. (2009): “An empirical investigation of the welfare effects of banning wholesale price discrimination,” *The RAND Journal of Economics*, 40(1), 20–46.

Yang, H., and L. Ye (2008): “Search with Learning: Understanding Asymmetric Price Adjustments,” *The RAND Journal of Economics*, 39(2), 547–564.

Yoshida, Y. (2000): “The Effects of Third-Degree Price Discrimination in Oligopoly,” *The American Economic Review*, 90(1), 244–246.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004): “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets,” *Review of Economics and Statistics*, 86(2), 514–533.

7 Appendix: Figures and tables

7.1 Figures

Figure 1: Markup variation, by product

Product A   Product B

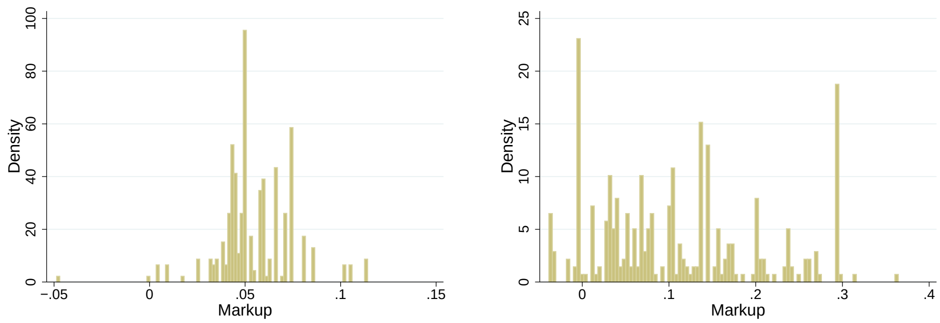

	Share of markup variation that is between-customer (i.e., cross-section variation), across products
Mean	.6468
Min	.2389
Max	.9535
St. Dev.	.1947

Notes: The figure and table were constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Markup is defined as (price-cost)/price. The between-customer variation of each product is computed using the between-within variance decomposition method of panel data.

Figure 2: Probability of breaking the relationship on price change, and on markup; at the transaction level

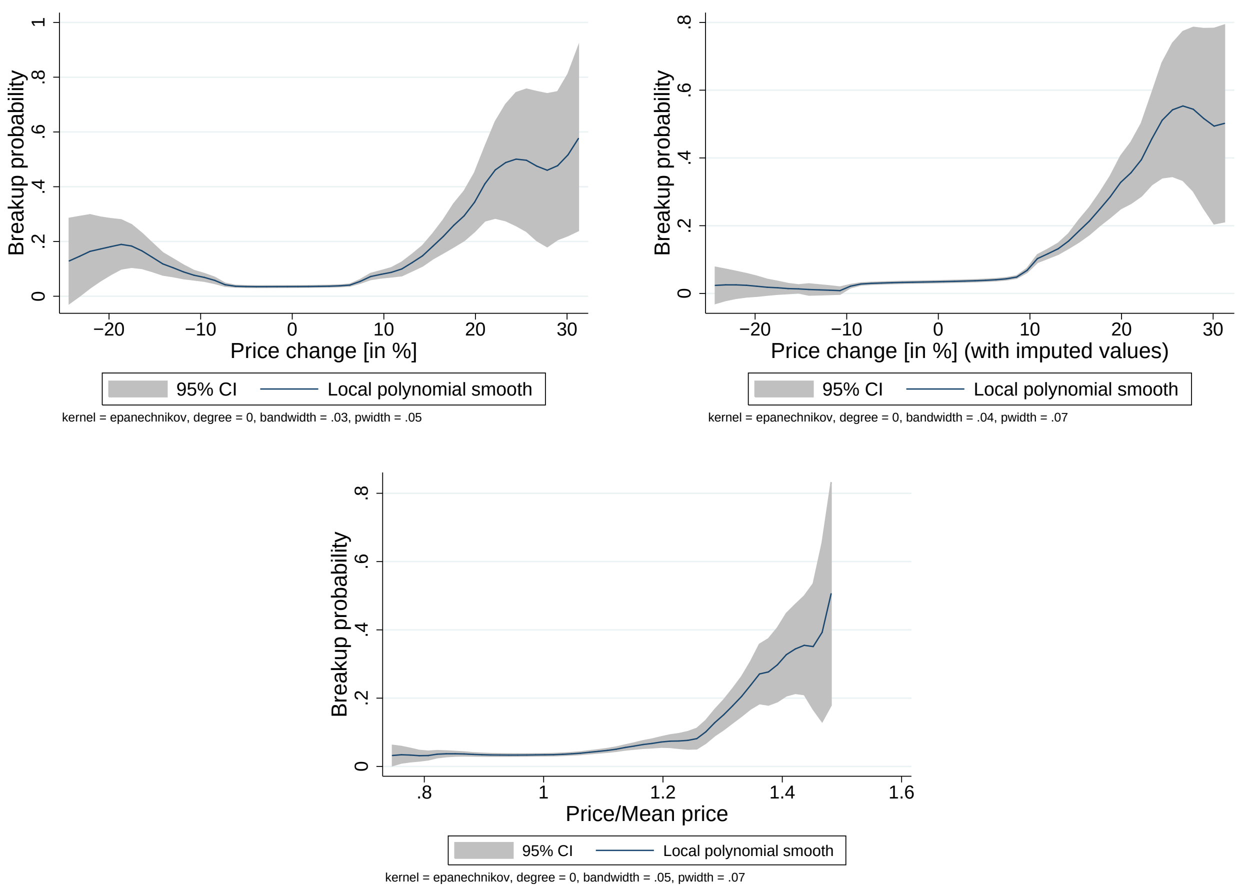

Notes: The figure was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. A breakup is customer-product specific, and is defined as a "longer than usual" interruption to the normal frequency of purchases. In most cases, the interruption extends beyond the sample period. Transactions near the end of the sample period are excluded. Price change [in %] is the percentual change in price relative to the price of the previous transaction. The price change [in %] variable with some imputed values is discussed in detail in Section 2.2. The price/mean price variable is the transaction price divided by the product-specific average price.

Figure 3: CDF of price change and markups, by breakup status.

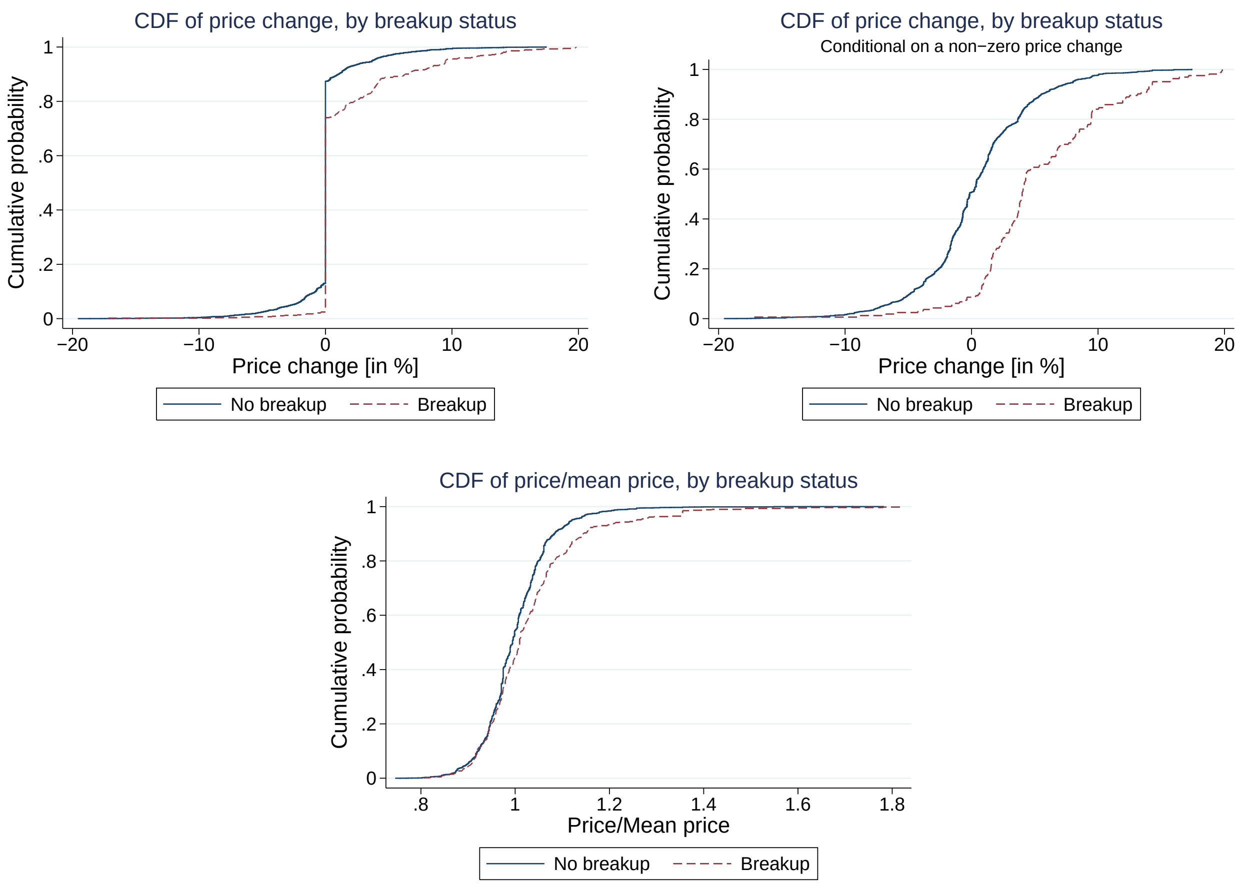

Notes: The figure was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. I construct the two top graphs using the price change [in %] variable with some imputed values, discussed in Section 2.2. The price/mean price variable is the transaction price divided by the product-specific average price.

Figure 4: Price change on cost change, at the transaction level: Price stickiness

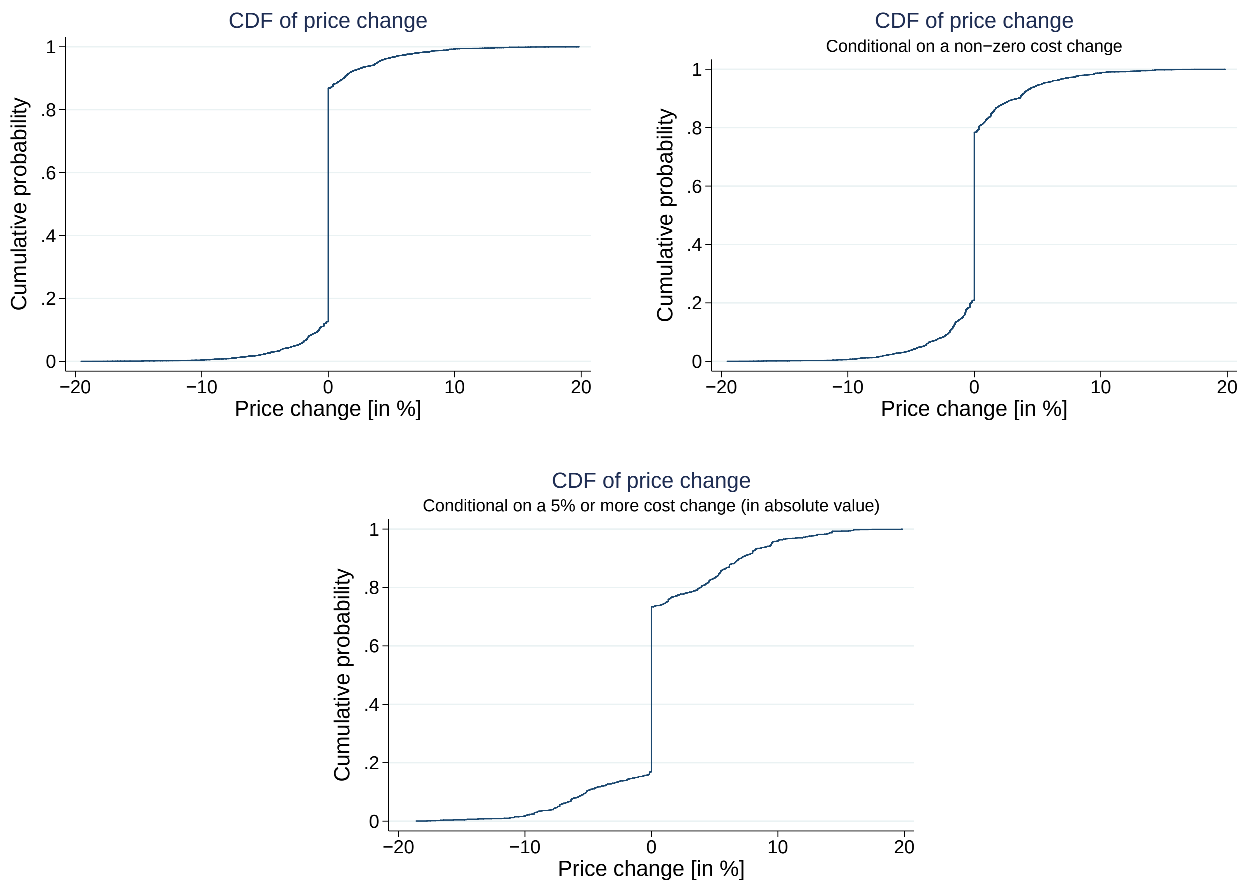

Notes: The figure was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Price change is defined as the event in which two consecutive transactions of a given customer-product combination display different prices. Cost change is defined analogously.

Figure 5: Price change decision on cost change, and markup deviation (defined as the deviation with respect to average customer-product markup), at the transaction level

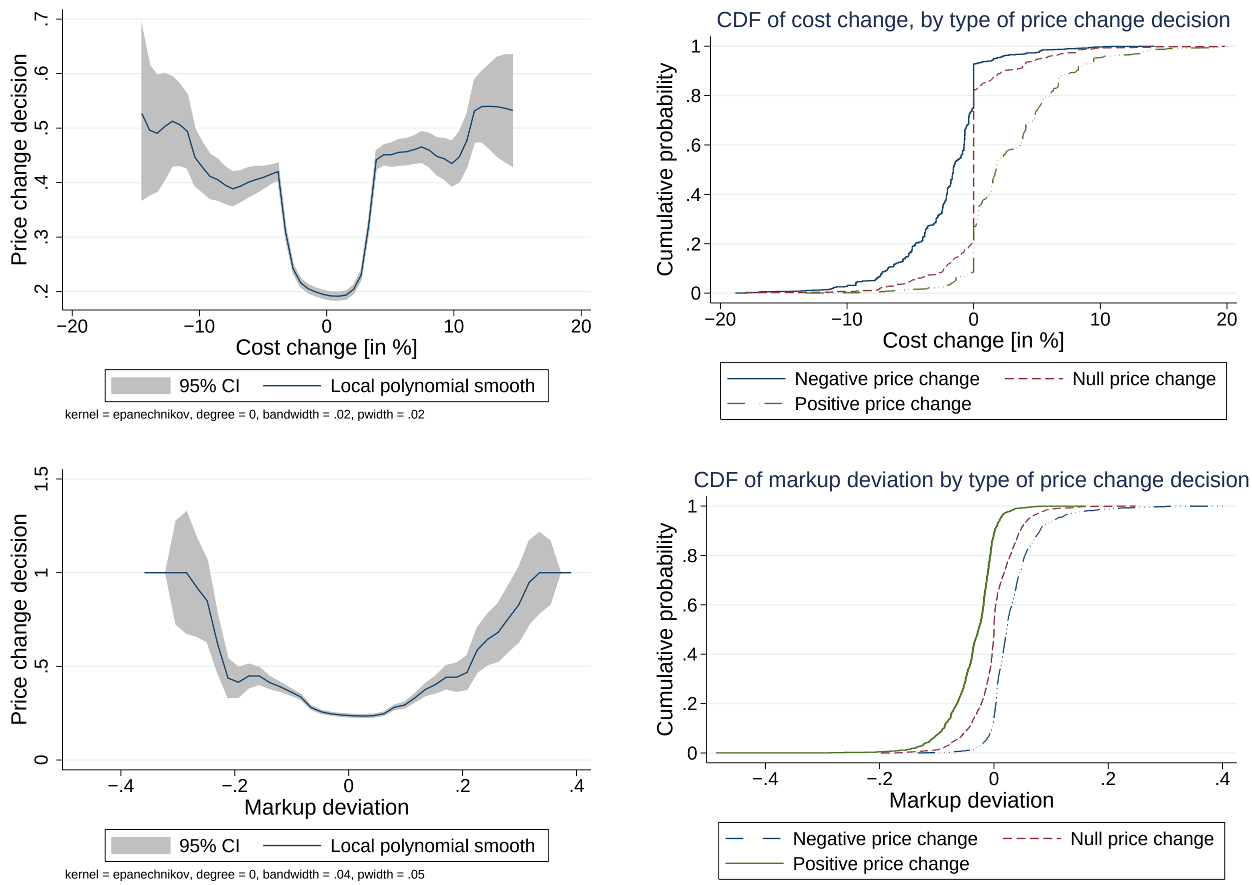

Notes: The figure was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Price change is defined as the event in which two consecutive transactions of a given customer-product combination display different prices. Markup deviation is given by the difference between the default markup before price change (i.e., using the price in the last transaction), and the average customer-product markup.

Figure 6: Price change value on markup deviation (defined as the deviation with respect to average customer-product markup), and cost change, at the transaction level

Notes: The figure was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. See footnote of Figure 5.

Figure 7: Timing of the game

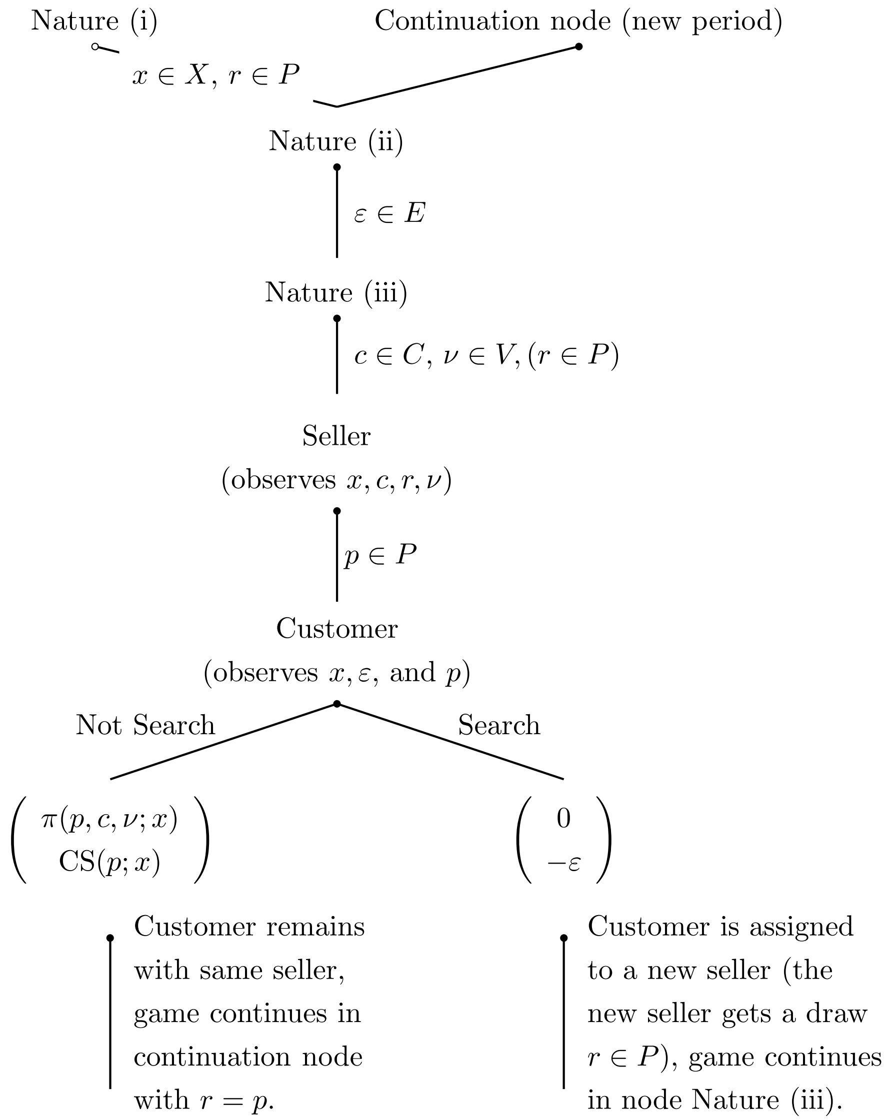

Figure 8: Comparative statics: Simulation results

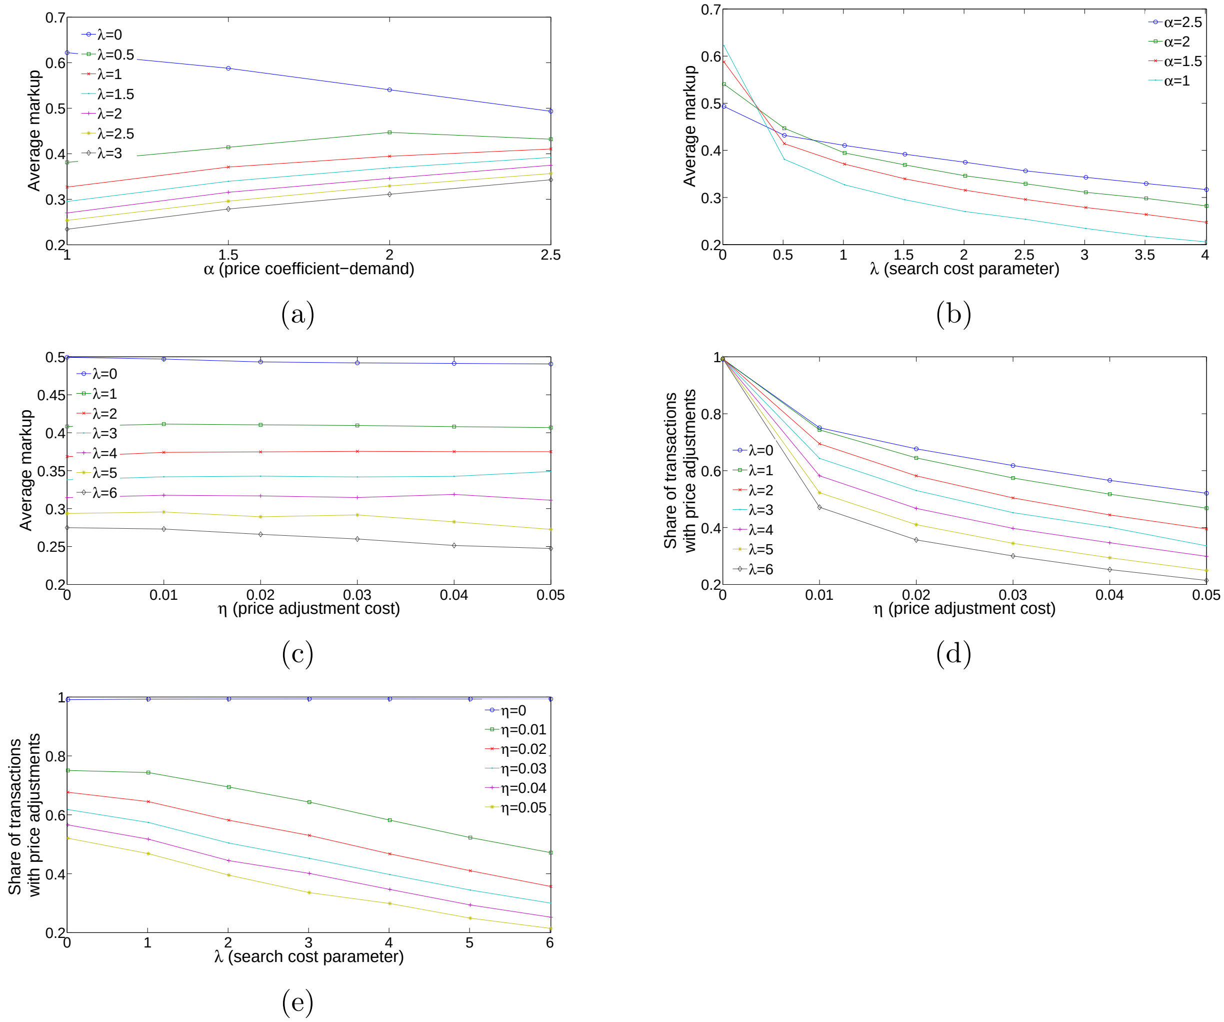

Findings: a) More price elastic consumers pay lower markups only when the price-elasticity of the probability of not searching is low enough. b) Customers with lower search costs (higher λ) pay lower markups. c) The price adjustment cost does not affect the average markup. d) The price adjustment cost (η) has a negative effect on the frequency of price adjustments. e) A higher value of the search cost parameter (λ) intensifies the effect of the price adjustment cost on the frequency of price changes but is not able to generate infrequent price changes itself.

Figure 9: Reservation price and search costs

(a) Low-markup product: Reservation prices

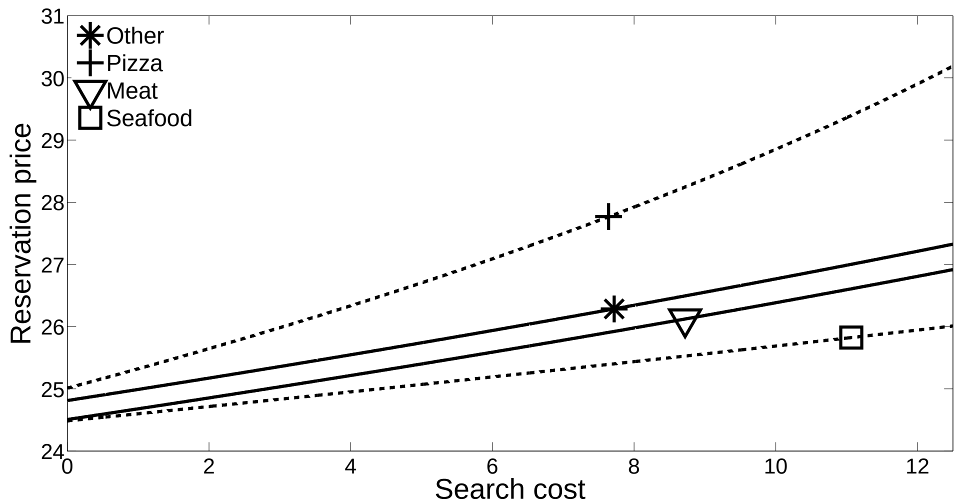

(b) High-markup product: Reservation prices

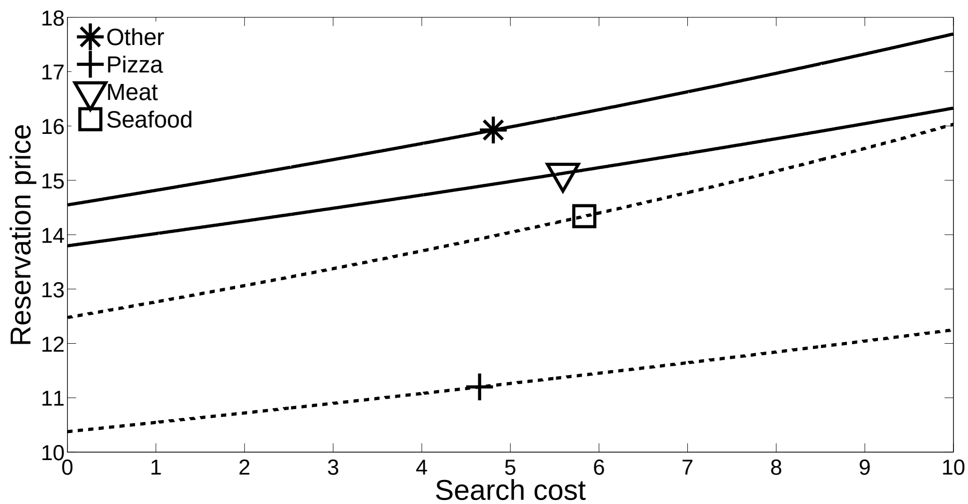

Notes: The graphs plot the reservation price as a function of search costs (see Section 3.2 for a formal discussion). The horizontal position of the marker on each curve indicates the mean search cost of the corresponding customer type.

7.2 Tables

Table 1: Cost changes in counterfactual transaction, by type of price change in transaction

	Mean cost change in counterfactual transaction	
Type of price change	No breakup	Breakup
Negative	-.00604424*	.03419901*
Null	.00275295	.00965499
Positive	.0045786	.00753719

* $p < 0.05$

Notes: The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. The counterfactual transaction is a fictitious transaction after the last record that I place in time using the frequency and time between purchases of each customer-product combination. The cost change is computed using the cost at the time of the counterfactual transaction (observed for transactions happening sufficiently far from the end of the sample period). The type of price change is the one that the customer-product combination saw in her last recorded transaction. See Subsection 2.2 for further discussion.

Table 2: Probability of breakup on markup: OLS regressions

	(1)	(2)	(3)	(4)
	Breakup of business relationship			
Price/Mean price	0.173*	0.130*	0.161*	0.120*
	(0.0349)	(0.0367)	(0.0348)	(0.0374)
Price/Mean price * 1{Price change > 0}			0.0223*	0.0102+
			(0.00601)	(0.00521)
Customer FE	No	Yes	No	Yes
Product FE	No	Yes	No	Yes
Observations	15576	15576	15576	15576
R^2	0.005	0.242	0.006	0.242

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$, + $p < 0.1$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. The price/mean price variable is the transaction price divided by the product-specific average price. A breakup is customer-product specific, and is defined as a "longer than usual" interruption to the normal frequency of purchases. In most cases, the interruption extends beyond the sample period. Transactions near the end of the sample period are excluded.

Table 3: Probability of breakup on price change [in percentage]: OLS regressions

	(1)	(2)	(3)
	Breakup of business relationship		
Price change [in %]	0.00580*	0.00450*	0.00325*
	(0.000994)	(0.000861)	(0.000816)
Price change [in %, Lag]		0.000652	0.000543
		(0.000830)	(0.000807)
Price change [in %, Lag 2]			0.000136
			(0.000831)
Customer FE	Yes	Yes	Yes
Product FE	Yes	Yes	Yes
Observations	15507	14341	13346
R^2	0.234	0.202	0.193

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. I construct this table using the price change [in proportion] variable with some imputed values, discussed in Subsection 2.2. See the footnote of Table 2 for details on the construction of the breakup variable.

Table 4: Number of breakups, by customer-time combination.

Item	Number	Per cent
0	7,484	94
1	438	5
2	62	1
3	11	0
4	1	0
5	1	0
7	1	0
Total	7,998	100

The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer-time combination.

Table 5: Explaining variation of the number of unique products purchased per month: OLS regressions

	(1)	(2)	(3)
	Number of unique products		
Customer FE	No	Yes	Yes
Restaurant type-Month FE	Yes	No	Yes
Observations	1807	1807	1807
R^2	0.037	0.874	0.878

The table was constructed using all of the observations in the data. An observation is the number of unique products purchased by a customer at a given month.

Table 6: Explaining the variation in markups: OLS regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Markup						
Product FE	Yes	No	No	No	No	Yes	No
Customer FE	No	No	No	No	No	Yes	No
Restaurant type-Product FE	No	Yes	No	No	Yes	No	No
Zip code-Product FE	No	No	Yes	No	Yes	No	No
Salesmen-Product FE	No	No	No	Yes	Yes	No	No
Customer-Product FE	No	No	No	No	No	No	Yes
Observations	15576	15576	15054	15576	15054	15576	15576
R^2	0.437	0.471	0.755	0.645	0.803	0.660	0.828

Notes: The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Markup is defined as (price-cost)/price.

Table 7: Markups on the percentile rank of the transaction quantity: OLS regressions

	(1) Markup	(2) Markup
Percentile rank of quantity	-0.0318*	-0.0261*
	(0.00873)	(0.00497)
Customer FE	No	Yes
Product FE	No	Yes
Observations	15576	15576
R^2	0.011	0.664

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. The percentile rank is the position of the transaction quantity in the CDF of quantity of the corresponding product.

Table 8: Markups and physical distance: OLS regressions

	(1)	(2)	(3)	(4)
	Markup			
Log(Minutes)		0.000517		
		(0.00401)		
Log(Miles)			-0.000938	
			(0.00260)	
Product FE	Yes	Yes	Yes	Yes
Salesman FE	Yes	Yes	Yes	Yes
Restaurant type FE	Yes	Yes	Yes	Yes
Zip code FE	No	No	No	Yes
Observations	15576	14960	14960	15054
R^2	0.463	0.472	0.472	0.537

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Markup is defined as (price-cost)/price. The distance, both in miles and minutes, was computed using Google Maps from the warehouse of the food distributor to the centroid of the zip code where each restaurant is located.

Table 9: Markups on a proxy for search costs: OLS regressions

	(1)	(2)
	Markup	
Share of total expenditure	-0.0758*	-0.0198*
	(0.00909)	(0.00803)
Percentile rank of quantity	-0.0140	-0.0188*
	(0.00833)	(0.00522)
Customer FE	No	Yes
Product FE	No	Yes
Observations	15576	15576
R^2	0.090	0.665

Notes: Clustered standard errors (at the customer-product level) in parantheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. The share of total expenditure variable is the share of total expenditure over time spent on that given product. One would expect that the higher the relative importance of a product for a customer's budget, the higher the willingness to search for better deals elsewhere. Percentile rank of quantity is defined in Table 7.

Table 10: Markups on a proxy for search costs (II): OLS regressions

	(1)	(2)
	Markup	
Other products in the same order/	0.0214*	0.00667*
Products ever bought	(0.0100)	(0.00321)
Customer FE	No	Yes
Product FE	No	Yes
Observations	15576	15576
R^2	0.003	0.661

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Other products in the same order/Products ever bought is the number of others products purchased in the same order (e.g. if it was a single-product order, the number of other products equals zero), as a fraction of the number of products that were ever purchased by that customer. The higher the fraction, the more likely it is that the customer has fewer providers, increasing his search costs.

Table 11: Are salesmen forward looking? Probability of price change on future cost changes: OLS regressions

	(1)	(2)	(3)	(4)
	Positive Price Change when ΔCost [t]>5%		Negative Price Change when ΔCost [t]<-5%	
1{ΔCost [t+1]>5%}	0.178*	0.205*		
	(0.0462)	(0.0547)		
1{ΔCost [t+1]<-5%}			0.0805	-0.0944
			(0.0808)	(0.110)
Customer FE	No	Yes	No	Yes
Product FE	No	Yes	No	Yes
Observations	1097	1097	818	818
R^2	0.014	0.579	0.001	0.721

Notes: Clustered standard errors (at the customer-product level) in parentheses. * $p < 0.05$. The table was constructed using a subset of 23 products that i) have at least 250 recorded transactions in the data, and ii) present cost variation over time. An observation is a customer–product–time combination. Columns 1 and 2 (3 and 4) are restricted to the observations for which there was at least a 5 percent cost increase (decrease). 1{ΔCost [t+1]>5%} is a dummy that takes the value of 1 if in the following transaction of the corresponding customer-product combination, there was at least a 5 percent cost increase.

Table 12: Customer Demand: Linear Regressions

	Low-markup product		High-markup product	
	OLS	2SLS	OLS	2SLS
Price	-0.0834*	-0.104*	-0.0637*	-0.112*
	(0.0112)	(0.0138)	(0.0294)	(0.0403)
Price * Pizza	-0.0174*	-0.0209*	0.0162	-0.00141 *
	(0.00126)	(0.00150)	(0.00834)	(0.00826)
Price * Meat	0.000868	0.000314	0.0112*	0.00665*
	(0.00191)	(0.00214)	(0.00303)	(0.00316)
Price * Seafood	0.0189*	0.0165*	-0.0189*	-0.0235*
	(0.00174)	(0.00198)	(0.00345)	(0.00378)
Observations	2244	2186	1478	1421
R^2	0.192	0.192	0.040	0.033

Robust standard errors in parentheses. * $p < 0.05$.

Notes: The table was constructed using the observations of two products, which are actively transacted, and differ in their average markups. An observation is a customer–product–time combination. The instrumental variables for the 2SLS regression include the wholesale cost, and zip code, and salesmen fixed effects. The F-statistic of the first-stage regressions of all price variables are above 100.

Table 13: Pseudo-Maximum Likelihood Estimates

	Low-markup product		High-markup product	
Parameter:	Estimate	St. error	Estimate	St. error
log(Adj. Cost)	-3.7355	0.1814	-2.1596	0.0191
$\mu_{1,0}$	2.0437	0.0344	1.5702	0.0107
μ_{Pizza}	-0.0103	0.0001	-0.0327	0.0002
μ_{Meat}	0.1219	0.0014	0.1513	0.0012
μ_{Fish}	0.3601	0.0072	0.1935	0.0014
$\tilde{\sigma}_\nu$	0.2939	0.0029	0.0908	0.0009
γ_{Other}	4.0511	0.0148	3.5692	0.0013
γ_{Pizza}	3.322	0.0015	4.0085	0.0447
γ_{Meat}	4.0899	0.0871	3.8064	0.0031
γ_{Seafood}	4.6633	0.0052	3.3165	0.0547
Pseudo Log Likelihood	-18,019.64		-3,568.64	
N	2,233		1,473	

Notes: The table was constructed using the observations of two products, which are actively transacted, and differ in their average markups. As discussed in Section 4, the exponential of parameter Adj. Cost determines the price adjustment cost, the parameter vector μ determines the means of the search cost distributions, the exponential of $\tilde{\sigma}_\nu$ determines the standard deviation of the distribution of unobserved shocks to the wholesale cost σ_ν, and finally, the exponential of the γ parameters give the equilibrium search reward. Standard errors were computed using the bootstrap, with $n = 25$ bootstrap samples.

Table 14: Price elasticity computed at the customer-type average price

	Low-markup product	High-markup product
Customer Type	Elasticity	Elasticity
Other	2.83	1.76
Pizza	3.46	1.80
Meat	2.85	1.70
Seafood	2.39	2.34
Weighted average	2.94	1.82

Notes: Price elasticities are computed using the estimates in Table 12, and the average price paid by each customer type in equilibrium. I report the absolute values.

Table 15: Average price elasticity of probability of serving the customer at the customer-type average price

Customer Type	Low-markup product Elasticity	High-markup product Elasticity
Other	11.53	7.34
Pizza	6.14	10.86
Meat	10.71	8.53
Seafood	13.72	6.13
Weighted average	10.35	8.18

Notes: Price elasticities are computed using the estimates in Table 13, and the type–dependent equilibrium price distribution. I report the absolute values.

Table 16: Contrasting rate of price change in the data with the estimated model predictions, by customer type

	Low-markup product		High-markup product	
Customer Type	Data	Model	Data	Model
Other	0.5595	0.8620	0.2963	0.3077
Pizza	0.7406	0.8424	0.0417	0.7575
Meat	0.7252	0.8507	0.7476	0.1290
Seafood	0.8592	0.8719	0.2316	0.5096

Notes: The table was constructed using the observations of two products, which are actively transacted, and differ in their average markups. Rate of price change is the share of transactions in which a price change is observed, where price change is defined as the event in which two consecutive transactions of a given customer-product combination display different prices.

Table 17: Welfare analysis: Comparing current market structure with perfect competition and the monopoly case (results reported as a share of total surplus under perfect competition).

Low-markup product

	Monopoly		Current		
Customer Type	Profits	TS	CS	Profits	TS
Other	0.1948	0.3896	0.5766	0.0630	0.6396
Pizza	0.1719	0.3437	0.4777	0.0796	0.5573
Meat	0.1951	0.3901	0.5867	0.0716	0.6583
Seafood	0.2152	0.4304	0.6476	0.0614	0.7090
Overall	0.1951	0.3901	0.5757	0.0658	0.6415

High-markup product

	Monopoly		Current		
Customer Type	Profits	TS	CS	Profits	TS
Other	0.1762	0.3523	0.4906	0.0493	0.5399
Pizza	0.1746	0.3491	0.6118	0.0836	0.6954
Meat	0.1840	0.3679	0.5410	0.0595	0.6005
Seafood	0.1513	0.3025	0.4918	0.1007	0.5924
Overall	0.1748	0.3496	0.5254	0.0622	0.5876

Notes: The welfare measures are the expected customer surplus and profits that are obtained in a completed transaction for the different market structures. These measures reflect the outcome of the search procedure (in the case of the current market structure). I report the welfare measures as a share of the total surplus under perfect competition.

Table 18: By what factor does the wholesale cost distribution need to change to justify the merger?

Customer Type	Low-markup product	High-markup product
Other	0.82	0.75
Pizza	0.86	0.61
Meat	0.81	0.70
Seafood	0.79	0.68
Overall	0.82	0.70

Notes: This table reports the factor by which the wholesale cost distribution has to be scaled with to make the monopoly total surplus match the total surplus under the current market structure. See Table 17 for details on the welfare measures.

Table 19: Welfare implications of shifting from the uniform-price equilibrium (U) to the price discrimination equilibrium (PD), by customer type

Low-markup product

Customer Type	$\Delta\mathbb{E}$[Search reward]	$\Delta\mathbb{E}$[CS]	$\Delta\mathbb{E}$[Profits]	$\Delta\mathbb{E}$[TS]
Other	1.62%	0.70%	-5.22%	0.09%
Pizza	-3.21%	-2.56%	14.78%	-0.41%
Meat	-0.55%	0.62%	4.33%	1.01%
Seafood	1.13%	1.42%	-11.81%	0.12%
Total change	0.76%	0.44%	-2.66%	0.12%

High-markup product

Customer Type	$\Delta\mathbb{E}$[Search reward]	$\Delta\mathbb{E}$[CS]	$\Delta\mathbb{E}$[Profits]	$\Delta\mathbb{E}$[TS]
Other	3.30%	2.99%	-18.03%	0.63%
Pizza	1.01%	0.21%	7763.56%	13.71%
Meat	-0.13%	0.77%	59.99%	4.60%
Seafood	-10.82%	-2.55%	512.48%	13.70%
Total change	1.58%	1.59%	54.05%	5.39%

Notes: These are changes in the expected value of each welfare measure (as discussed in Section 5), computed using the equilibrium joint distribution of price and both components of the wholesale cost under both the price discrimination case (PD) and the uniform pricing case (U). The welfare measures are the expected customer surplus and profits that are obtained in a completed transaction. These measures reflect the outcome of the search procedure. $\mathbb{E}$[Search reward] is the expected customer surplus given the equilibrium price distribution, recall from Section 3.5 that an increase in $\mathbb{E}$[Search reward] implies higher competition, and lower profits.