

Relative Price Dispersion: Evidence and Theory*

GREG KAPLAN

University of Chicago and NBER

LEENA RUDANKO

FRB Philadelphia

GUIDO MENZIO

University of Pennsylvania and NBER

NICHOLAS TRACHTER

FRB Richmond

August 2016

Abstract

We use a large dataset on retail pricing to document that a sizeable portion of the cross-sectional variation in the price at which the same good trades in the same period and in the same market is due to the fact that stores that are, on average, equally expensive set persistently different prices for the same good. We refer to this phenomenon as relative price dispersion. We show that relative price dispersion might stem from sellers' attempts to discriminate between high-valuation buyers who need to make all of their purchases in the same store and low-valuation buyers who are willing to purchase different items from different stores. We calibrate our theory and show that it is not only consistent with the extent and sources of dispersion in the price that different sellers charge for the same good, but also with the extent and sources of dispersion in the prices that different households pay for the same basket of goods.

JEL Codes: L11, D40, D83, E31.

Keywords: Price Dispersion, Equilibrium Product Market Search.

*We are grateful to Bob Hall for his comments on the previous work of some of the authors, which inspired the current project. We are grateful to Jim Albrecht, George Alessandria, Fernando Alvarez, Ken Burdett, Jeff Campbell, Glenn Ellison, Sarah Ellison, Doireann Fitzgerald, Gaetano Gaballo, Matthew Gentzkow, Erik Hurst, Nir Jaimovich, Boyan Jovanovic, Jose Moraga-Gonzales, Stephan Seiler, Rob Shimer, and Randy Wright for useful comments. We are grateful to our audiences at the NBER Summer Institute (IO and Price Dynamics groups); Minnesota Macro Workshop; Canadian Macro Study Group; Philadelphia Search and Matching Workshop; Econometric Society; Chicago FRB/St.Louis FRB Workshop on Money, Banking, Payments and Finance; SED Annual Meeting; Bank of France Price Setting and Inflation Conference; EIEF Junior Conference in Macroeconomics; University of Pennsylvania (Wharton); University of Wisconsin at Madison; FRB Minneapolis; George Washington University; Universidad Torcuato di Tella; Universidad San Andres; and UVA/Richmond FRB Research Jamboree. We thank the Kilts Center for Marketing and the Nielsen Company for sharing their data. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Philadelphia, Federal Reserve Bank of Richmond, or the Federal Reserve System.

We document that a significant fraction of the cross-sectional variation in the price at which the same good is sold in the same period of time and in the same market is due to the fact that retailers that are, on average, equally expensive set persistently different prices for that particular good. We refer to this phenomenon as *relative price dispersion*. We develop a theory of relative price dispersion in an equilibrium multiproduct model of the retail market. According to our theory, relative price dispersion stems from the sellers' incentive to discriminate between high-valuation buyers who need to make all of their purchases in the same store, and low-valuation buyers who are willing to purchase different items from different stores. We calibrate our theory and show that it can simultaneously capture the extent and sources of dispersion in the price that different sellers charge for the same good, as well as the extent and sources of dispersion in the prices that different households pay for the same basket of goods.

In the first part of the paper, we measure the extent and sources of dispersion in the price at which the same good is sold in the same week and in the same geographical area by different retailers. We carry out the analysis using the Kilts-Nielsen Retail Scanner (KNRS) dataset, which provides weekly price and quantity information for 1.5 million goods (defined by their universal product code or UPC) at around 40,000 stores in more than 2,500 counties across 205 designated market areas (DMA), which are geographical areas of roughly the same size as metropolitan statistical areas.

Using the KNRS, we compute the average price of a particular good in a given week and in a given geographical area. We normalize the price of the good at each different store by expressing it as a percentage deviation from the average price in the relevant week and market. We break down the normalized price of the good at each different store into a *store component*—defined as the average of the normalized price of all the goods sold by the store in the week—and a *store-good component*—defined residually as the difference between the price of the good at the store and the average price of the store. We then compute the variance of the price of the good at different stores, and we break down this variance into a store component and a store-good component. We find that, on average, the standard deviation of prices for the same good in the same week and market is 15.3%, and the variance is 2.34%. We also find that only 15% of the variance is due to the variance in the store component—i.e., due to the fact that the same good is sold at stores that have a different average price—while 85% is due to the variance of the store-good component—i.e., due to the fact that the same good is sold at different prices at stores that are equally expensive on average.

We then use the time dimension in KNRS to identify the *persistent component* and the

transitory component of the price of a particular good at a particular store. To this aim, we follow an approach that is commonly used in the literature on labor economics to analyze wage inequality (see, e.g, Gottschalk and Moffitt 1994 and Blundell and Preston 1998) but that had never been applied before to study price dispersion. Specifically, we estimate a statistical model for prices, in which both the store and the store-good component of prices are given by the sum of a fixed effect, an AR process and an MA process. When estimated, the statistical model fits the auto-covariances of prices very well. The estimated model implies that almost all of the cross-sectional variance of the store component of prices stems from persistent differences in average store prices. Also, the estimated model implies that roughly one-third of the variance of the cross-sectional variance of the store-good component of prices stems from persistent differences in the price of the good at equally expensive stores, while roughly two-thirds of the variance stems from transitory differences. Overall, a sizable fraction of the variation in the price of the same good across different stores in the same week and in the same market is due to the fact that equally expensive stores charge *persistently different* prices for the same good. We refer to the persistent differences in the price of a particular good relative to the average price of the store as relative price dispersion.

The literature offers compelling theories of transitory differences in the price of the same good across equally expensive stores. For instance, according to the theory of intertemporal price discrimination (see, e.g., Conlisk, Gerstner and Sobel 1984; Sobel 1984; and Menzio and Trachter 2015a), sellers find it optimal to occasionally lower the price of a particular good in order to discriminate between low-valuation customers who are willing to do their shopping at any time during the month, and high-valuation customers who need to make their purchases on a specific day of the month. As different sellers implement these occasional price reductions at different times, the equilibrium may feature short-term differences in the price of the same good across equally expensive stores. According to the inventory management theory (see, e.g., Aguirregabiria 1999), a seller finds it optimal to increase the price of a good as the inventory of the good falls, and to lower the price when the inventory of the good is replenished. As different sellers have different inventory cycles, the equilibrium may feature short-term differences in the price of the same good across equally expensive stores. In contrast, the literature does not offer much theoretical guidance in the way of understanding relative price dispersion, i.e., long-term differences in the price of the same good across equally expensive stores.

In the second part of the paper, we develop a theory of relative price dispersion. We consider a retail market in which sellers and buyers respectively supply and demand two goods. Sellers are ex-ante homogeneous, both with respect to their cost of producing the goods and in terms of the type of buyers they meet. Buyers are ex-ante heterogeneous. One

type of buyers, which we call *busy*, has a relatively high valuation for goods and needs to make all purchases in the same store. The other type of buyers, which we call *cool*, has a relatively low valuation and can purchase different items at different stores. The retail market is imperfectly competitive. As in Butters (1977) and Burdett and Judd (1983), we assume that buyers do not have access to all the sellers but only to a subset of them. In particular, a fraction of buyers can access a single seller, and a fraction of buyers can access multiple sellers.

We find that, for some parameter values, the equilibrium of the retail market must display relative price dispersion. This result follows from two properties of the equilibrium: (i) some sellers find it optimal to post different prices for the two goods; and (ii) for every seller that posts a lower price for one good, there is another seller posting a lower price for the other good. There is a simple intuition behind the first property of the equilibrium. Consider a seller that sets the same price for the two goods and suppose that this common price lies in between the valuation of the cool buyers and the valuation of the busy buyers. Given this common price, the seller trades with some of the busy buyers but never trades with the cool buyers. Now, suppose that the seller lowers the price of the first good and increases the price of the second good, so as to keep the average price constant. Since the busy buyers must purchase both goods in the same location, they are indifferent to the change in prices. Hence, the seller trades the same quantity of goods to this type of buyer. However, as the seller keeps lowering the price of the first good, at some point it reaches the valuation of the cool buyers. Hence, at some point, the seller trades the good to this type of buyer. Overall, the seller is strictly better off setting different prices for the two goods, rather than setting a common price. The second property of the equilibrium follows from the fact that, if there were more sellers charging a lower price for one good than for the other, then there would be some unexploited profit opportunities. When taken together, the two properties of the equilibrium imply that there is dispersion in the price of the same good among sellers that are, on average, equally expensive.

According to our theory, relative price dispersion is the equilibrium consequence of a new type of price discrimination. The difference in valuation between the busy buyers and the cool buyers gives sellers a reason to try to price discriminate. The difference in the ability of the busy and the cool buyers to make their purchases at different locations gives the seller a way to price discriminate. Specifically, price discrimination is achieved by setting asymmetric prices for the two goods, so as to charge a high price for the basket of goods to the high-valuation buyers who need to purchase all items from the same retailer, and to charge a low price (on one item) to the low-valuation buyers who have the ability to purchase different items from different retailers.

In the last part of the paper, we calibrate and validate our theory of relative price dispersion. We calibrate the theory so as to match the extent of dispersion in the persistent component of prices at which different stores sell the same good, the contribution of the (persistent) store component and of the (persistent) store-good component to the dispersion of prices, and the elasticity of the price paid by a household for a given basket of goods with respect to the number of stores from which the household shops. We find that our theory can match these features of the data very well. Dispersion in the (persistent) store component of prices is obtained for the same reason as in the single product retail market models of Butters (1977) and Burdett and Judd (1983). Dispersion in the (persistent) store-good component of prices is an equilibrium consequence of price discrimination between low-valuation buyers who can shop at multiple locations and high-valuation buyers who must shop at a single location. The negative elasticity of the price paid by a household for a given basket of goods with respect to the number of stores from which it shops is obtained because, when there is relative price dispersion, buyers can achieve lower prices by purchasing different items at different stores.

We validate the theory using the Kilts-Nielsen Consumer Panel (KNCP), which tracks the shopping behavior of 50,000 households drawn from 54 different markets. We compute the extent and sources of dispersion in the prices that different households pay for the same basket of goods in the same quarter and in the same market. After removing the dispersion due to transitory variation in the store-good component of prices, we find that the standard deviation of prices paid by different households for the same basket of goods is 7.8%, and the variance is 0.61%. We also find that 55% of the variance is due to persistent differences in the store component of prices, while 45% is due to persistent differences in the store-good component of prices (and to the covariance term). We then compute the extent and the sources of dispersion in the prices paid by different households that are implied by our theory. We find that the predictions of the theory line up very well with the data. The theory correctly predicts that the dispersion of prices paid by different households for the same basket of goods is lower than the dispersion of prices posted by different stores for the same good. Also, the theory correctly predicts that variation in the store component of prices contributes to a larger share of the dispersion of prices paid by households than to the dispersion of prices posted by stores. Intuitively, the predictions of the theory are correct because the variation of the store-good component of prices of individual goods washes out in the basket price of those households who need to purchase all goods from the same store.

Overall, our model of the retail market simultaneously matches reasonably well three features of the data: (i) the variance of the price charged by sellers for the same good in the same week and in the same market, and the magnitude of the different components of the

variance; (ii) the variance of the price paid by households to purchase the same basket of goods in the same quarter and in the same market, and the magnitude of the different components of the variance; (iii) the elasticity of the price paid by a household for a given basket of goods with respect to the number of stores from which the household shops. We believe that this is an important accomplishment. Any compelling theory of price dispersion should not only generate the variance of prices observed in the data, but it should also generate this variance through the same channels as in the data. Moreover, any compelling theory of price dispersion should square well with the extent and sources of dispersion in the prices paid by different households for the same basket of goods. After all, most theories of pricing in the retail market are fundamentally theories of the shopping behavior of buyers (see, e.g., Varian 1980, Burdett and Judd 1983, Conlisk, Gerstner and Sobel 1983). Ours is the first theory of price dispersion to be confronted—and to match reasonably well—these additional features of the data and, hence, we believe it to help us better understand the retail market. In turn, we believe that understanding the retail market is a useful pursuit. For economists in Marketing, understanding the retail market is key to assess the profitability of sellers’ pricing strategies. For economists in Industrial Organization, understanding the retail market is critical to evaluate regulatory interventions. For economists in Macro, understanding the retail market and, especially, the shopping behavior of households is useful to measure individual consumption (see, e.g., Aguiar and Hurst 2005), to measure inequality (see Atanasio and Pistaferri 2016), and even to understand business cycles (see, e.g., Kaplan and Menzio 2016).

The paper contributes to the large empirical literature documenting price dispersion. Ours is the first paper to use a (large-scale) dataset that covers multiple products, each sold at multiple stores and each observed over a long period of time. Only this type of data allows estimation of the overall variance of prices, the variance due to temporary and persistent differences in the store component of prices, and the variance due to temporary and persistent differences in the store-good component of prices. The variance decomposition identifies a novel feature of retail prices: relative price dispersion. Part of the previous empirical literature has documented the extent of price dispersion for particular products: cars and anthracite coal in Stigler (1961), 39 products in Pratt, Wise, and Zeckhauser (1979), several books and CDs in Brynjolfsson and Smith (2000), four academic textbooks in Hong and Shum (2006), gasoline in Lewis (2006), illegal drugs in Galenianos, Pacula, and Persico (2012), mortgage brokerage services in Woodward and Hall (2012), and many more. There are also studies that have used small-scale data that covers multiple products, each sold at multiple stores and each observed over time: car insurance policies in Alberta in Dahlby and West (1986), prescription drugs in upstate New York in Sorensen (2000), and four products

in Israel in Lach (2002). However, none of these papers has attempted to decompose price dispersion as in our paper and, hence, identified relative price dispersion.¹

The paper also contributes to Search Theory. In particular, we develop a version of the search-theoretic model of equilibrium price dispersion of Burdett and Judd (1983) in which buyers and sellers trade multiple goods and buyers are heterogeneous with respect to both their valuation and their ability to purchase different items at different stores. Our version of the model leads to relative price dispersion and fits well both data on the dispersion of prices across stores and data on the dispersion of prices across households. There are many search-theoretic models of equilibrium price dispersion in the retail market for a single good: Butters (1977), Varian (1980), Burdett and Judd (1983), and Stahl (1989) in which price dispersion emerges because some buyers are in contact with a single seller and others are in contact with multiple sellers; Reinganum (1979) and Albrecht and Axell (1984) in which price dispersion emerges because of heterogeneity among buyers or sellers; Rob (1985), Stiglitz (1987), and Menzio and Trachter (2015b) in which price dispersion is obtained because sellers are large and have an impact on reservation prices. There are fewer search-theoretic models of equilibrium price dispersion in the retail market for multiple goods, as these models are famously difficult to analyze. The main exceptions are McAfee (1995), Zhou (2014), Baughman and Burdett (2015), Rhodes (2015), and Rhodes and Zhou (2015). Yet, none of these papers considers the type of buyers' heterogeneity on which our theory of relative price dispersion is built.

In Industrial Organization, there are several Hotelling-style models of the retail market in which a seller finds it optimal to charge different prices for goods that are intrinsically identical (see, e.g., Lal and Matutes 1994). Asymmetric pricing is optimal because it is assumed that buyers are aware of the price of only a subset of goods when deciding where to shop. There are also some versions of these models in which, as in ours, buyers differ with

¹A predecessor of our paper is Kaplan and Menzio (2015). Kaplan and Menzio (2015) decompose the price of a good in a transaction into three components: (i) a store component (defined as the store-level price), (ii) a store-good component (defined as the average price of the good during a quarter relative to the store-level price), (iii) a transaction component (defined residually). They then break down price dispersion into variance of the store component, variance of the store-good component and variance of the transaction component. Their price dispersion decomposition is similar in spirit to ours, where the store-good component is meant to capture the persistent part of the store-good component of the price and the transaction component is meant to capture the transitory part of the store-good component. The problem with Kaplan and Menzio (2015) is that they use a much smaller dataset than ours and, for many store-good combinations, they have very few observations per quarter. Therefore, the estimates of the permanent and transitory parts of the store-good component of the price are very noisy. So much so that, in most of the paper, Kaplan and Menzio (2015) only consider their sum. A successor of our paper is Gorodnichenko, Sheremirov, and Talavera (2016). The focus of their paper is on the properties of online prices. Among many other findings, they show that relative price dispersion is an important contributor to overall price dispersion also in online markets. See also Nakamura and Steinsson (2008) and Campbell and Eden (2014).

respect to their ability to purchase different items in different locations (see, e.g., Lal and Matutes 1989, and Chen and Rey 2012). These models, even though different from ours in many dimensions, can generate relative price dispersion under some conditions. Unlike ours, these models have not been confronted with the data.²

1 Relative Price Dispersion: Evidence

In this section, we jointly analyze the dispersion and dynamics of the prices of identical goods in the same geographical area and over the same period of time. We use a detailed dataset on prices that includes the time series of the price of a large number of goods at each of a large number of stores. We use these data to estimate a rich stochastic process for the average price level of a store and for the price of a good at a store relative to the average price level of the store. We then use the estimated stochastic process to decompose the variance of the price of the same good in the same period of time and same geographic area. The new finding that we emphasize is that a substantial fraction of the cross-sectional variance of prices is due to the fact that stores that are, on average, equally expensive set persistently different prices for the same good. We refer to this phenomenon as relative price dispersion.


1.1 Framework and Estimation Strategy

Let p_{jst} denote the quantity-weighted average price of good j at store s in time period t . In our application, a time period is one week and a good is defined by its UPC (barcode). We first decompose the log of each price p_{jst} into three additively separable components: a component that reflects the average price of the good in period t , μ_{jt} ; a component that reflects the expensiveness of the store selling the good, y_{st} ; and a component that reflects factors that are unique to the combination of store and good, z_{jst} .³ Formally, we decompose

²There is also a literature in Marketing that frames the retailer’s pricing decision as a choice between Every Day Low Pricing (EDLP) and Promotional pricing (PROMO). EDLP is a pricing strategy involving selling always at “low” prices. PROMO is a strategy involving deep, temporary sales on a subset of goods. Part of this literature (see, e.g., Bell and Lattin 1998) studies the preferences of different types of buyers for stores that follow different pricing strategies. Another part of this literature (see, e.g., Ellickson and Misra 2008) studies the retailers’ choice between these two pricing strategies. As far as we know, this literature does not offer an explanation for relative price dispersion. That is, it does not answer the question “Why do sellers that are, on average, equally expensive choose to set persistently different prices for the same good?” If sellers follow PROMO strategies that are asynchronized, there will be equally expensive stores with different prices for the same good. However, these differences should be transitory, not persistent. And, even if the temporary sales were very long-lasting, why should they be asynchronized across stores?

³We work with the natural logarithm of quantity-weighted average prices. This reflects an assumption that innovations to prices enter multiplicatively, which is convenient when jointly analyzing prices of many

Figure 1: Auto-correlation function of prices



Notes: The figure plots the empirical auto-correlation functions of the store and store-good components, \hat{y}_{st} and \hat{z}_{jst} , together with their counterparts from the fitted statistical model.

the log of p_{jst} as

$$\log p_{jst} = \mu_{jt} + y_{st} + z_{jst}. \quad (1)$$

We model both the store component of the price, y_{st} , and the store-good component of the price, z_{jst} , as the sum of a fixed effect, a persistent part and a transitory part. This statistical model is motivated by the empirical shape of the auto-correlation functions of y_{st} and z_{jst} , which are illustrated in Figure 1. The auto-correlation functions of y_{st} and z_{jst} display a sharp drop at short lags, followed by a smoothly declining profile that remains strictly positive even at very long lags. The initial drop in the auto-correlation suggests the presence of a transitory component in both y_{st} and z_{jst} . We model the transitory components as an MA(q) process, rather than an IID process, to allow for the possibility that the transitory component may reflect temporary sales. Indeed, since sales may last longer than one week and since the timing of sales may not correspond to the weekly reporting periods, they are better captured by a process with some limited persistence than with a weekly IID process.⁴ The smoothly declining portion of the auto-correlation function is consistent with the presence of an AR(1) component. Finally, the fact that the auto-correlation function remains positive even after 100 weeks suggests the presence of a fixed effect.

Motivated by the above observations, we use the following statistical model for y_{st} and different goods.

⁴In Appendix A, we show that our findings are robust to alternative specifications for the process of the transitory part of the store-good component of prices.

z_{jst} :

$$\begin{aligned}
y_{st} &= y_s^F + y_{st}^P + y_{st}^T, & z_{jst} &= z_{js}^F + z_{jst}^P + z_{jst}^T, \\
y_{st}^P &= \rho_y y_{s,t-1}^P + \eta_{s,t}^y, & z_{jst}^P &= \rho_z z_{js,t-1}^P + \eta_{jst}^z, \\
y_{st}^T &= \varepsilon_{s,t}^y + \sum_{i=1}^q \theta_{y,i} \varepsilon_{s,t-i}^y, & z_{jst}^T &= \varepsilon_{jst}^z + \sum_{i=1}^q \theta_{z,i} \varepsilon_{jst-i}^z, \\
y_s^F &= \alpha_s^y, & z_{js}^F &= \alpha_{js}^z,
\end{aligned} \tag{2}$$

where y_s^F and z_{js}^F denote the fixed effects of the store and the store-good components, y_{st}^P and z_{jst}^P denote the persistent parts of the store and the store-good components, and y_{st}^T and z_{jst}^T denote the transitory parts of the store and the store-good components. The parameters α_s^y and α_{js}^z are random variables with mean zero and variance $\sigma_{\alpha^y}^2$ and $\sigma_{\alpha^z}^2$. The parameters ρ_y and ρ_z are the auto-regressive parameters of the AR(1) part of the store and store-good components, while $\eta_{s,t}^y$ and η_{jst}^z are the innovations to the AR(1) part and are assumed to be random variables with mean zero and variance $\sigma_{\eta^y}^2$ and $\sigma_{\eta^z}^2$. Finally, the parameters $\theta_{y,i}$ and $\theta_{z,i}$ are the coefficients of the MA(q) part of the store and store-good components, while $\varepsilon_{s,t}^y$ and ε_{jst}^z are the innovations to the MA(q) part and are assumed to be normal random variables with mean zero and variance $\sigma_{\varepsilon^y}^2$ and $\sigma_{\varepsilon^z}^2$. All random variables are independent across goods, stores, and times. In our baseline model, we set $q = 1$.

We estimate the parameters of the statistical model in (2) using prices p_{jst} for a large number of goods $j = 1 \dots J$, at a large number of stores $s = 1 \dots S$, for a long sequence of weeks $t = 1, \dots, T$, in a single geographic market m . Given the large number of goods, stores, and time periods, and the presence of unobserved components in prices, estimating this model via maximum likelihood (or with panel data instrumental variables regressions) is not feasible. Instead, we estimate the model using a multistage generalized method of moments approach that is analogous to techniques that are commonly used when estimating models of labor earnings dynamics (see, e.g., Gottschalk and Moffitt 1994, Blundell and Preston 1998, Kaplan 2012).

The estimation procedure involves four steps.

Step 1. We estimate the good-time mean, μ_{jt} , as the average of the log price, $\log p_{jst}$, across all stores s in the market of interest, i.e.,

$$\hat{\mu}_{jt} = \frac{1}{S} \sum_{s=1}^S \log p_{jst}. \tag{3}$$

We construct normalized prices as

$$\tilde{p}_{jst} = \log p_{jst} - \hat{\mu}_{jt}. \tag{4}$$

Step 2. We estimate the store component, \hat{y}_{st} , by taking sample means of the normalized prices across all goods in store s , i.e.,

$$\hat{y}_{st} = \frac{1}{n_{jst}} \sum_{j=1}^{n_{jst}} \tilde{p}_{jst}, \quad (5)$$

where n_{jst} is the number of goods for which we have data for store s in period t . In some instances, $n_{jst} < J$ because not every store-good combination will meet our sample selection requirements in every week. We estimate the store-good component, z_{jst} , as

$$\hat{z}_{jst} = \tilde{p}_{jst} - \hat{y}_{st}. \quad (6)$$

The above process leads to a $S \times T$ panel of store components $\{\hat{y}_{st}\}$ and a $(J \times S) \times T$ panel of store-good components $\{\hat{z}_{jst}\}$ (where there may be missing data for some combinations of (j, s, t)).

Step 3. We construct the auto-covariance matrix of each of these panels up to L lags.

Step 4. We minimize the distance between the theoretical auto-covariance matrices implied by the model and the empirical auto-covariance function from step three. We use a diagonal weighting matrix that weights each moment by $n_{jst}^{0.5}$. However, the main results are not sensitive to using an identity weighting matrix instead.

1.2 Kilts-Nielsen Retail Scanner Dataset

We estimate the statistical model in (2) using the KNRS dataset. The KNRS dataset contains store-level weekly sales and unit average price data at the UPC level. The dataset covers the period 2006 to 2012. The full dataset contains weekly price and quantity information for more than 1.5 million UPCs at around 40,000 stores in more than 2,500 counties across 205 DMAs. A DMA is a geographic area defined by Nielsen that is roughly the same size as a metropolitan statistical area. Since our estimation procedure requires computing a full auto-covariance matrix at the store-good-week level, it is not feasible to estimate the model using anywhere near the full set of UPCs. For example, in the Minneapolis-St Paul DMA alone, the full data set would consist of more than 200 million observations of p_{jst} per year. Thus, in order to keep the size of the analysis manageable, we restrict attention to a subset of the data.

We focus our analysis on a single DMA: Minneapolis-St Paul. However, there is nothing particular about Minneapolis-St-Paul. Indeed, in Appendix A, we show that our findings are robust to the choice of different markets (i.e., different DMAs) and to different definitions

of markets (i.e., Counties or States rather than DMAs). For the Minneapolis-St Paul DMA, we focus on the 1,000 UPCs with the largest quantities of sales in the state of Minnesota during the first quarter of 2010. Table 9 in Appendix B shows how these 1000 products are distributed across the 10 departments defined by Nielsen. Table 8 in Appendix B shows how these products are distributed across 50 of the 125 product groups defined by Nielsen. Among these 1000 products, the one with the most units sold belongs to the “Fresh Eggs” product module (2.9 million units). The one with the fewest units sold belongs to the “Liquid Cocktail Mixes” product module (50 thousand units).⁵ Even after restricting attention to these 1,000 products, the dataset is extremely large. Over the seven year period from 2006 to 2012, we have more than 40 million observations of prices p_{jst} . To ensure that our findings are not specific to this particular bundle of goods, in Appendix A we re-estimate the model using a number of alternative sets of UPCs, chosen in various ways.

Given our set of UPCs and our geographical area, we select stores, goods, and weeks that satisfy two criteria:

1. For each store/week combination, we have quantity and price data for at least N_1 of the UPCs in in the given set. In our baseline estimation below, we set $N_1 = 250$. In Appendix A, we report results for $N_1 \in \{50, 500\}$.
2. For each good/week combination, we have quantity and price data for at least N_2 stores. In our baseline estimation below, we set $N_2 = 50$. In Appendix A, we report results for $N_2 \in \{25, 100\}$.

These selection criteria ensure that we focus only on store/goods/weeks where we have sufficient data to reliably estimate the good-time means and store-time means in the first and second stages of the estimation procedure. In addition, to avoid the influence of large outliers when computing the empirical auto-covariance function, we drop observations of the store components and store-good components whose absolute value is greater than one.

1.3 Estimation and Variance Decomposition

Table 1 presents the results of our benchmark estimation of the statistical model. Figure 1 shows that the estimated statistical model fits very well the empirical auto-correlation function of both the store component of prices and of the store-good component of prices.

⁵Nielsen divides the full set of UPCs in the product database into 10 “departments,” which are subdivided into around 125 “product groups,” that are further subdivided into around 1,075 “product modules.” For example, different sized bottles of Heinz Tomato Ketchup have distinct UPCs in the “Catsup” product module, which is one of 34 product modules in the “Condiments, Gravies and Sauces” product group, one of 38 product groups in the “Dry Grocery” department.

Table 1: Parameter estimates

Store component	ρ^y	$\theta_{y,1}$	$Var(\alpha^y)$	$Var(\eta^y)$	$Var(\varepsilon^y)$
	0.983	0	0.00279	2.46E-05	0.000117
Store-good component	ρ^z	$\theta_{z,1}$	$Var(\alpha^z)$	$Var(\eta^z)$	$Var(\varepsilon^z)$
	0.965	0.0256	0.00327	0.000262	0.0127

Notes: The model is estimated on a baseline sample of UPCs using data for the Minneapolis-St Paul designated market area.

Let us explain how the estimated model achieves such a remarkable fit of the data. First, consider the store component of prices. The small drop in the auto-correlation at a lag of one is captured by the model with a relatively small standard deviation in the MA(1) part of the store component. The slow decay in the auto-correlation at lags greater than one is captured by the model by having an auto-regressive coefficient in the AR(1) part of the store component of 98.2% per week. Finally, the auto-correlation remains extremely high even at very long lags. Indeed, it is 80% even at a lag of 100 weeks. The model captures this feature of the data by having a standard deviation in the fixed effect of the store component of prices of 5.3%.

Next, consider the store-good component of prices. The large drop in the auto-correlation function at a lag of one is captured by the model with a relatively high standard deviation of the MA(1) part of the store-good component of the price. The slow decay in the auto-correlation at lags greater than one is captured by the model with an auto-regressive coefficient in the AR(1) part of the store-good component of 96.5% per week. Finally, the auto-correlation remains above 15% even at lags of about 2 years. The model captures this feature of the data by having a standard deviation in the fixed effect of the store-good component of prices of 5.5%. The model is only unable to capture the spike in the auto-correlation at 52 weeks (which, presumably, reflects annual patterns in the pricing behavior of stores) and the high-frequency zig-zagging of the auto-correlation (which reflects patterns in the behavior of sales that cannot be captured by our statistical model).

Since the estimated model fits the auto-correlation functions very well, we are comfortable using it to decompose the variance of the price at which the same good is sold by different stores in the same week and in the same market. Table 2 contains the results of the decomposition. We find that the overall variance of the price is 0.023, equivalent to a standard deviation of 15.3%. The variance of the store-component of the price accounts for only 15.5% of the overall variance, while the variance of the store-good component of the price accounts for the remaining 84.5%. In words, only a small fraction of price dispersion is due to the fact that the good is sold by stores that are, on average, more or less expensive. Most

Table 2: Dispersion in prices: persistent and transitory

	Variance	Percent		Std. Dev.
<u>Store component</u>				
Transitory	0.000	3.2		0.011
Fixed plus Pers.	0.004	96.8		0.059
Total Store	0.004	100.0	15.5	0.060
<u>Store-good component</u>				
Transitory	0.013	64.1		0.113
Fixed plus Pers.	0.007	35.9		0.084
Total Store-good	0.020	100.0	84.5	0.141
<u>Total</u>	0.023	100.0		0.153

Notes: The left panel presents the cross-sectional variances of UPC prices, as well as the store and store-good components separately. The middle panel presents the decomposition of this variance into persistent and transitory components. The right panel presents the cross-sectional standard deviations.

of price dispersion is due to the fact that stores that are, on average, equally expensive sell the same good at different prices. The result is consistent with earlier findings in Sorensen (2000), Campbell and Eden (2014) and Kaplan and Menzio (2016).

The variation in the store component of the price could be due to persistent or transitory differences in the store-level price. Similarly, the variation in the store-good component of the price could be due to persistent or transitory differences across stores in the price of the good relative to the store-level price. The statistical model (2) allows us to distinguish between these sources of variation. Since the estimated auto-regressive coefficient in the AR(1) part of both the store and the store-good component of the price is very high, we choose to define as permanent differences those associated with the fixed-effect and the AR(1). We define as transitory differences those associated with the MA(1).

Given the above definitions, we find that 96.8% of the variance in the store component of the price is due to persistent differences in the store-level price, while only 3.2% is due to transitory differences. Furthermore, we find that 35.9% of the variance in the store-good component is due to persistent differences across stores in the price of the good relative to the store-level price, while 64.1% is due to transitory differences. In words, nearly all of the differences in the expensiveness of different stores are persistent. In contrast, 2/3 of the differences across stores in the price of the good relative to the store-level price are transitory and 1/3 are persistent. The decomposition of the variance of the store and store-good component of the price into permanent and transitory parts is novel to this paper. Kaplan and Menzio (2015) did attempt such decomposition but, since they used a much smaller dataset, they were not able to draw reliable conclusions.

Variance decompositions are a convenient tool for breaking down dispersion into orthogonal elements. However, since variances are measured in squared prices, the interpretation of the magnitude of the various elements is somewhat hard to interpret. For this reason, the last column in Table 2 reports the standard deviation of each of the elements in our decomposition. We find that the standard deviation of the store component of the price is 6%. The standard deviation of the transitory part of the store-good component of the price is 11.3%. And the standard deviation of the persistent part of the store-good component of the price is 8.4%.

Our decomposition clearly reveals that a sizeable fraction of price dispersion is due to the fact that stores that are, on average, equally expensive set persistently different prices for the same good. This feature of the data, which we believe is first documented here, is what we call relative price dispersion. Relative price dispersion implies that, among equally expensive stores, some stores charge systematically a high price on some goods and a low price on other goods, while other stores do the opposite. Relative price dispersion seems, at first blush, hard to explain. After all, why would similar stores choose systematically different pricing strategies?

Before proposing an explanation for relative price dispersion, we want to reassure the reader that this feature of the data is very robust. In Appendix A, we show that relative price dispersion is still there (and is still quantitatively important) when we consider different inclusion criteria for the data, when we look at different markets, when we define markets more or less broadly, and when we model the transitory part of the store-good component of prices in a way that more closely resembles the typical pattern of temporary sales.

We also want to reassure the reader that there is no “banal” explanation for relative price dispersion. In Appendix A, we show that relative price dispersion does not seem to be caused by managerial inattention (see, e.g., Ellison, Snyder and Zhang 2015). Indeed, if similar stores set systematically different prices for the same good because of inattention, we would expect relative price dispersion to be only a feature of low-ticket items. In contrast, we find that, when we restrict attention to a subsample of more expensive goods, relative price dispersion is still sizeable (albeit smaller than for less expensive goods). We also show that relative price dispersion does not seem to be caused by differences in wholesaler-retailer relationships. Indeed, if similar stores charged systematically different prices for the same good because of better or worse relationships with wholesalers, relative price dispersion should disappear once we restrict attention to a subset of goods from a single wholesaler (as differences in prices across stores should then be absorbed by the store component of the price). In contrast, when we restrict attention to products from Unilever, we find that variance

decomposes in a way that is very similar to Table 2. We reach the same conclusions when we restrict attention to products from Coca-Cola. Finally, we show that relative price dispersion is not related to shelf management issues. We show that the variance decomposition is very similar for a subsample of goods with short shelf-life and for a subsample of goods with long shelf-life.

2 Relative Price Dispersion: Theory

In this section, we develop a theory of relative price dispersion. According to our theory, relative price dispersion emerges as a strategy that retailers use in order to price discriminate between high-valuation buyers who need to make all of their purchases in the same location, and low-valuation buyers who have the time to purchase different items in different places. In Section 2.1, we describe the model, which is an extension to multiple goods and multiple types of buyers of the canonical theory of price dispersion of Burdett and Judd (1983). In Section 2.2, we establish some general properties of equilibrium. In Section 2.3, we characterize the equilibrium when competition between sellers is weak. In this equilibrium, sellers post prices that are attractive only to the high-valuation buyers. Since high-valuation buyers purchase all goods in the same location, the model behaves like the one-good model of Burdett and Judd (1983). In Section 2.4, we characterize the equilibrium when competition between sellers is stronger. In this equilibrium, some sellers post some prices that are attractive to low-valuation buyers. Since low-valuation buyers can purchase different goods from different locations, the model does not behave like Burdett and Judd (1983), and, indeed, the model features relative price dispersion. In Section 2.5, we briefly describe the equilibrium when competition between sellers is strongest.

2.1 Environment

We consider a retail market populated by homogeneous sellers and heterogeneous buyers who trade two goods (i.e., good 1 and good 2). Specifically, the market is populated by a measure $s > 0$ of identical sellers.⁶ Every seller is able to produce each of the two goods at

⁶We assume that sellers are ex-ante identical. If sellers had different costs or faced different populations of buyers, it would be easy to generate relative price dispersion. However, such explanations of relative price dispersion would be basically unfalsifiable, as data on wholesale prices and demand curves faced by different retailers is generally unavailable. Indeed, as noted by Stigler (1961), “It would be metaphysical, and fruitless, to assert that all dispersion is due to heterogeneity.” To further strengthen this point, let us draw a parallel with the literature on temporary sales. Clearly, one could explain temporary sales with temporary declines in wholesale prices or with temporary increases in the elasticity of demand faced by retailers. Yet, because such explanations are hard to falsify and seem empirically implausible, the literature has developed theories in which sellers choose to have temporary sales in a stationary environment (see, e.g., Conlisk, Gerstner and

the same, constant marginal cost, which we normalize to zero. Every seller chooses a price for good 1, p_1 , and a price for good 2, p_2 , so as to maximize his profits, taking as given the distribution $H(p_1, p_2)$ of the vector of prices across sellers. We find it useful to denote as $F_i(p)$ the fraction of sellers whose price for good $i \in \{1, 2\}$ is smaller than p , and as $\lambda_i(p)$ the fraction of sellers whose price for good $i \in \{1, 2\}$ is equal to p . We refer to $F_i(p)$ as the distribution of prices for good $i \in \{1, 2\}$. Similarly, we find it useful to denote as $G(q)$ the fraction of sellers whose prices p_1 and p_2 sum up to less than q , and as $\nu(q)$ the fraction of sellers whose prices sum up to q . We refer to $G(q)$ as the distribution of basket prices.

On the other side of the retail market, there is a measure 1 of buyers. A fraction $\mu_b \in (0, 1)$ of buyers are of type b and a fraction $\mu_c = 1 - \mu_b$ of buyers are of type c , where b stands for *busy* and c stands for *cool*. A buyer of type b demands one unit of each good, for which he has valuation $u_b > 0$. A buyer of type c demands one unit of each good, for which he has valuation u_c , with $u_b > u_c > 0$. More specifically, if a buyer of type $i \in \{b, c\}$ purchases both goods at the prices p_1 and p_2 , he attains a utility of $2u_i - p_1 - p_2$. If a buyer of type $i \in \{b, c\}$ purchases one of the two goods at the price p , he attains a utility of $u_i - p$. If a buyer of type $i \in \{b, c\}$ does not purchase any of the goods, he attains a utility of zero.

In the retail market, trade is frictional. We assume that a buyer cannot purchase from just any seller in the market, as each buyer only has access to a small network of sellers. In particular, a buyer of type b can access only one seller with probability α_b and multiple (namely, two) sellers with probability $1 - \alpha_b$, with $\alpha_b \in (0, 1)$. Similarly, a buyer of type c can access only one seller with probability α_c and two sellers with probability $1 - \alpha_c$, with $\alpha_c \in (0, 1)$. For the sake of simplicity, we let $\alpha_b = \alpha_c = \alpha$.⁷ We refer to a buyer who can only access one seller as a *captive buyer*, and to a buyer who can access multiple sellers as a *non-captive buyer*. We interpret these restrictions on the buyers' access to sellers as physical constraints (i.e., sellers that the buyer can easily reach) rather than as informational constraints⁸ (i.e. sellers of which the buyer is aware). We also assume that a buyer of type b must always make all of his purchases from just one of the sellers in his network. In contrast, a buyer of type c can purchase different goods from different sellers in his network. We interpret this assumption as heterogeneity in the buyer's ability or willingness to visit multiple stores when shopping.

A few comments about the environment are in order. We consider a version of the canonical model of price dispersion of Butters (1977) and Burdett and Judd (1983). In this

Sobel 1984; Sobel 1984; or Albrecht, Postel-Vinay and Vroman 2013).

⁷It is tedious but straightforward to generalize our theory of relative price dispersion to the case of $\alpha_c > \alpha_b$.

⁸The reader may want to think of our model as a version of Hotelling (1929), in which buyers have a zero transportation cost to access some sellers and an infinite transportation cost to access some other sellers.

model, the co-existence of buyers who have access to only one seller and of buyers who have access to multiple sellers induces identical sellers to post different prices for the same good. We depart from the canonical model by considering a market in which buyers and sellers trade two goods. Obviously, to develop a theory of relative price dispersion, it is necessary to consider multiproduct retailing. The simplest case of multiproduct retailing involves two products. We also depart from the canonical model by considering a market in which buyers are heterogeneous. In particular, we assume that buyers of type b are willing to pay higher prices for the goods than buyers of type c and that they are less willing to purchase different goods from different retailers than buyers of type c . It is natural to think about type- b buyers as buyers whose time has high value in the labor market, and about type- c buyers as buyers whose time has low value in the labor market. Hence, type- b buyers are willing to pay higher prices, but they are hesitant to spend time shopping at different retailers.⁹ The negative correlation in the buyers' willingness to pay and ability to shop from multiple retailers is the key to our theory of relative price dispersion.

Our model of the retail market is static, as Butters (1977) and Burdett and Judd (1983). We interpret the equilibrium price distribution of our static model as a long-term outcome. Indeed, in a repeated version of the model, it can be seen immediately that sellers would have nothing to gain from changing their prices over time. Moreover, in the presence of any type of adjustment costs, sellers would face a loss from changing their prices over time. Thus, in a repeated version of the model, sellers would keep their prices constant. Given our interpretation of the model, we compare the equilibrium price distribution to the distribution of the persistent component of sellers' prices.¹⁰

2.2 General Properties of Equilibrium

We start the analysis of the model by establishing some general properties of equilibrium. First, we argue that there are no sellers whose prices (p_1, p_2) sum up to more than $2u_b$. Second, we consider sellers whose prices (p_1, p_2) sum up to strictly more than $u_b + u_c$. These are sellers who cannot price an individual good below the willingness to pay of type- c buyers without raising the price of the other good above the willingness to pay of type- b buyers. We show that, in any equilibrium, the prices of these sellers must lie in a particular region

⁹As suggested by Kaplan and Menzio (2016), the reader may want to think of type- b buyers as employed ones, and of type- c buyers as the unemployed and/or retirees.

¹⁰It would be easy to extend our model to capture the high-frequency fluctuations in sellers' prices. Menzio and Trachter (2015a) use a version of Butters (1977) and Burdett and Judd (1983) to show that, if the buyers who are more willing to intertemporally substitute their purchases are also in contact with more sellers (or have a lower willingness to pay for goods), then each seller will post different prices at different times of the week in order to discriminate between different types of buyers.

in \mathbb{R}_+^2 . Third, we consider sellers whose prices sum up to less than $u_b + u_c$ and strictly more than $2u_c$. These are sellers who can price one of the goods below the willingness to pay of type- c buyers, while keeping the price of the other good below the willingness to pay of type- b buyers. We show that, in any equilibrium, the prices of these sellers must lie in one of two particular regions in \mathbb{R}_+^2 .

Lemma 1 contains two results. First, the lemma shows that a seller never finds it optimal to post a price greater than the willingness to pay of type- b buyers. This result is intuitive, as neither buyers of type b nor buyers of type c are willing to purchase goods at a price greater than u_b . Second, the lemma establishes that, if the sum q of the prices posted by a seller is strictly greater than $u_b + u_c$, then the seller's prices p_1 and p_2 are both strictly greater than the willingness to pay of type- c buyers and smaller than the willingness to pay of type- b buyers. That is, if the sum q of the prices posted by a seller is strictly greater than $u_b + u_c$, the seller's prices p_1 and p_2 must fall in the R_1 region in Figure 2. This result is also intuitive. If a seller has a basket price strictly greater than $u_b + u_c$, he must post $p_1 > u_c$ and $p_2 > u_c$ in order to make sure that both of his prices are smaller than u_b .

Lemma 1: (i) In any equilibrium, no seller posts prices (p_1, p_2) , such that $p_1 > u_b$ and/or $p_2 > u_b$. (ii) In any equilibrium, a seller with a basket price $q > u_b + u_c$ posts prices $(p_1, p_2) \in R_1$, where R_1 is defined as


$$R_1 = \{(p_1, p_2) : p_1 \in (u_c, u_b], p_2 \in (u_c, u_b], p_1 + p_2 > u_c + u_b\}. \quad (7)$$

Proof: Appendix C. ■

Lemma 2 states that, if the sum q of the prices posted by a seller is smaller than $u_b + u_c$ (but strictly greater than $2u_c$), then the seller's prices p_1 and p_2 are such that one of them is smaller than the willingness to pay of type- c buyers, and the other one is greater than the willingness to pay of type- c buyers. That is, if the sum q of the prices posted by a seller is smaller than $u_b + u_c$ and strictly greater than $2u_c$, the seller's prices p_1 and p_2 must fall in one of the two regions marked R_2 in Figure 2.

Let us provide some intuition for Lemma 2. Consider a seller with a basket price of $q \in (2u_c, u_b + u_c]$ who posts the same price $p_1 = p_2 = q/2$ for both goods. The seller will trade with some type- b buyers. However, the seller will not trade with any type- c buyers, as the price of each good is above their willingness to pay, u_c . Now suppose that the seller starts lowering the price of one good and, at the same, starts increasing the price of the other good in order to keep the price of the basket constant. The seller will trade with the same number of type- b buyers as before because these buyers, having to make all of their purchases in the same place, only care about the price of the basket. However, when the

Figure 2: Pricing decision of sellers



Notes: This figure considers the pricing decision of sellers discussed in the text, illustrating which regions of the (p_1, p_2) space will not be profit-maximizing. Conditional on a basket price in the interval $u_b + u_c \leq q \leq 2u_b$, sellers will not price outside R_1 . Conditional on a basket price in the interval $2u_c \leq q \leq u_b + u_c$, sellers will not price outside R_2 .

seller brings the price of the cheaper good below u_c , he will also start to trade this good to some type- c buyers. Therefore, a seller with a basket price of $q \in (2u_c, u_b + u_c]$ will never find it optimal to post the same price for both goods. For this seller, it is optimal to price the two goods asymmetrically.

Lemma 2: *In any equilibrium, a seller with a basket price $q \in (2u_c, u_b + u_c]$ posts prices $(p_1, p_2) \in R_2$, where R_2 is defined as*

$$R_2 = \{(p_1, p_2) : p_1 \in [0, u_c], p_2 \in (u_c, u_b], p_1 + p_2 \in (2u_c, u_b + u_c]\} \cup \{(p_1, p_2) : p_2 \in [0, u_c], p_1 \in (u_c, u_b], p_1 + p_2 \in (2u_c, u_b + u_c]\}. \quad (8)$$

Proof: Suppose that there is an equilibrium where a seller posts prices (p_1, p_2) , with $q \equiv p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in (u_c, u_b)$ and $p_2 \in (u_c, u_b)$. The seller attains a profit of

$$S(p_1, p_2) = \mu_b [\alpha + 2(1 - \alpha)(1 - G(q) + \nu(q)/2)] q. \quad (9)$$

Let us explain (9) in detail. The seller is in the network of $\mu_b \alpha$ captive buyers of type b . A captive buyer of type b will purchase both goods from the seller, as $p_1 \leq u_b$ and $p_2 \leq u_b$. The seller is in the network of $\mu_b 2(1 - \alpha)$ non-captive buyers of type b . If the second retailer in the network of such a buyer has a basket price $q' > q$, the buyer will purchase both goods from the seller. In fact, if the buyer purchases both goods from the seller, he attains a utility of $2u_b - q$. If the buyer purchases both goods from the second retailer, he attains

a utility of $2u_b - q'$ (recall that all sellers post prices below u_b). If the buyer purchases only one good, he attains a utility that is strictly smaller than $u_b - u_c$, which is smaller than $2u_b - q$. If the second retailer in the network of the buyer has a basket price $q' = q$, the buyer is indifferent between purchasing both goods from the seller or from the second retailer. We assume that, in this case, the buyer will randomize. If the second retailer in the network of the buyer has a basket price $q' < q$, the buyer will not purchase anything from the seller. Overall, a non-captive buyer of type b will purchase both goods from the seller with probability $1 - G(q) + \nu(q)/2$, and will purchase nothing from the seller with complementary probability. The seller is also in the network of $\mu_c\alpha$ captive buyers of type c , and of $\mu_c2(1 - \alpha)$ non-captive buyers of type c . However, a buyer of type c does not purchase any good from the seller, as both p_1 and p_2 are strictly greater than u_c .

If the seller deviates from the equilibrium and posts prices (p'_1, p'_2) with $p'_1 = u_c$, $p'_2 = q - p'_1$, he attains a profit of

$$\begin{aligned} S(p'_1, p'_2) &= \mu_b [\alpha + 2(1 - \alpha)(1 - G(q) + \nu(q)/2)] q \\ &\quad + \mu_c [\alpha + 2(1 - \alpha)(1 - F_1(p'_1) + \lambda_1(p'_1)/2)] p'_1. \end{aligned} \tag{10}$$

The expression for $S(p'_1, p'_2)$ is easy to understand. The seller is in the network of $\mu_b\alpha$ captive buyers of type b , and he will sell both goods to each of them. The seller is in the network of $\mu_b2(1 - \alpha)$ non-captive buyers of type b , and he will sell both goods to a fraction $1 - G(q) + \nu(q)/2$ of them. The seller is in the network of $\mu_c\alpha$ captive buyers of type c , and he will sell good 1 to each of them. Finally, the seller is in the network of $\mu_c2(1 - \alpha)$ non-captive buyers of type c . If the second retailer in the network of such a buyer has a price for good 1 greater than p'_1 , the buyer will purchase good 1 from the seller. If the second retailer in the network of such a buyer has a price for good 1 equal to p'_1 , the buyer will randomize. If the second retailer in the network of such a buyer has a price for good 1 smaller than p'_1 , the buyer will purchase good 1 from the second retailer. Overall, a non-captive buyer of type c will purchase good 1 from the seller with probability $1 - F_1(p'_1) + \lambda_1(p'_1)/2$.

Notice that the first term on the right-hand side of (10) is the same as $S(p_1, p_2)$. This is because buyers of type b base their purchasing decisions only on the price of the whole basket of goods, not on the price of individual goods. The second term on the right-hand side of (10) is strictly positive as it is bounded below by $\mu_c\alpha p'_1$. Overall, the right-hand side of (10) is strictly greater than $S(p_1, p_2)$. Therefore, there is no equilibrium in which a seller with a basket price of $q \in (2u_c, u_b + u_c]$ finds it optimal to post $p_1 \in (u_c, u_b)$ and $p_2 \in (u_c, u_b)$. This observation, combined with part (i) of Lemma 1, implies that, if a seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, he must be setting one of the two prices below u_c , and the other price in between u_c and u_b . ■

In Lemma 3, we establish two additional results. First, we show that the distribution of basket prices across sellers, $G(q)$, does not have any mass points. This property of equilibrium is obtained for the same reason as in Butters (1977), Varian (1980) or Burdett and Judd (1983). Specifically, if there is a mass point at $q_0 > 0$, a seller with a basket price of q_0 could lower one of his two prices by an arbitrarily small amount, and, instead of selling to half of the non-captive buyers of type b who are in touch with another seller charging q_0 , he could sell to all of them. Moreover, if there is a mass point at $q_0 = 0$, a seller with a basket price of q_0 could raise his prices to $p_1 = p_2 = u_b$ and, instead of attaining a profit of zero, he could attain a strictly positive profit.

Second, we show that the distribution of prices for individual good, $F_i(p)$, does not have any mass points for $p \in (0, u_c]$. The logic behind this result is also similar to the one in Butters (1977), Varian (1980), and Burdett and Judd (1983). However, in this case, we cannot rule out the possibility of a mass point at $p_0 > u_c$ or at $p_0 = 0$. There may be a mass point at $p_0 > u_c$ because, when the price of an individual good is higher than the willingness to pay of type- c buyers, the seller never trades the good in isolation and his price does not play any allocative role. There may be a mass point at $p_0 = 0$ because the fact that a seller trades one good at a price of zero does not imply that the seller's profit is zero.

Lemma 3: (i) In any equilibrium, $G(q)$ does not have any mass points. That is, $\nu(q) = 0$ for all q . (ii) In any equilibrium, $F_1(p)$ and $F_2(p)$ do not have any mass points for $p \in (0, u_c]$. That is, $\lambda_1(p) = 0$ and $\lambda_2(p) = 0$ for all $p \in (0, u_c]$.

Proof: Appendix C. ■

2.3 Bundled Equilibrium

In this section, we consider an equilibrium in which every seller in the market sets a basket price q strictly greater than $u_b + u_c$. We have already established that sellers with a basket price q strictly greater than $u_b + u_c$ post prices for the individual goods that are strictly greater than the willingness to pay of type- c buyers, and smaller than the willingness to pay of type- b buyers. Hence, sellers with a basket price q strictly greater than $u_b + u_c$ sell the whole bundle of goods to buyers of type b , and do not sell anything to buyers of type c . For this reason, we refer to this type of equilibrium as a Bundled Equilibrium. In the first part of this section, we characterize the price distribution in a Bundled Equilibrium. In the second part, we identify necessary and sufficient conditions for the existence of this type of equilibrium. Even though a Bundled Equilibrium is equivalent to the equilibrium of the one-good model of Burdett and Judd (1983) and does not necessarily feature relative price dispersion, the analysis in this section is useful as a stepping stone in the characterization

of more complex equilibria.

In a Bundled Equilibrium, all sellers set a basket price q greater than $u_b + u_c$. In Lemma 1, we showed that a seller with a basket price q greater than $u_b + u_c$ posts prices p_1 and p_2 that are strictly greater than u_c and smaller than u_b . That is, he posts prices $(p_1, p_2) \in R_1$. In Lemma 4, we show that a seller who posts a pair of prices $(p_1, p_2) \in R_1$ would attain the same profit by posting any other pair of prices $(p'_1, p'_2) \in R_1$ as long as $p'_1 + p'_2 = p_1 + p_2$.

Lemma 4: *The seller attains the same profit by posting the prices $(p_1, p_2) \in R_1$ and the prices $(p'_1, p'_2) \in R_1$ as long as $p'_1 + p'_2 = p_1 + p_2$.*

Proof: The profit of a seller posting $(p_1, p_2) \in R_1$ is given by

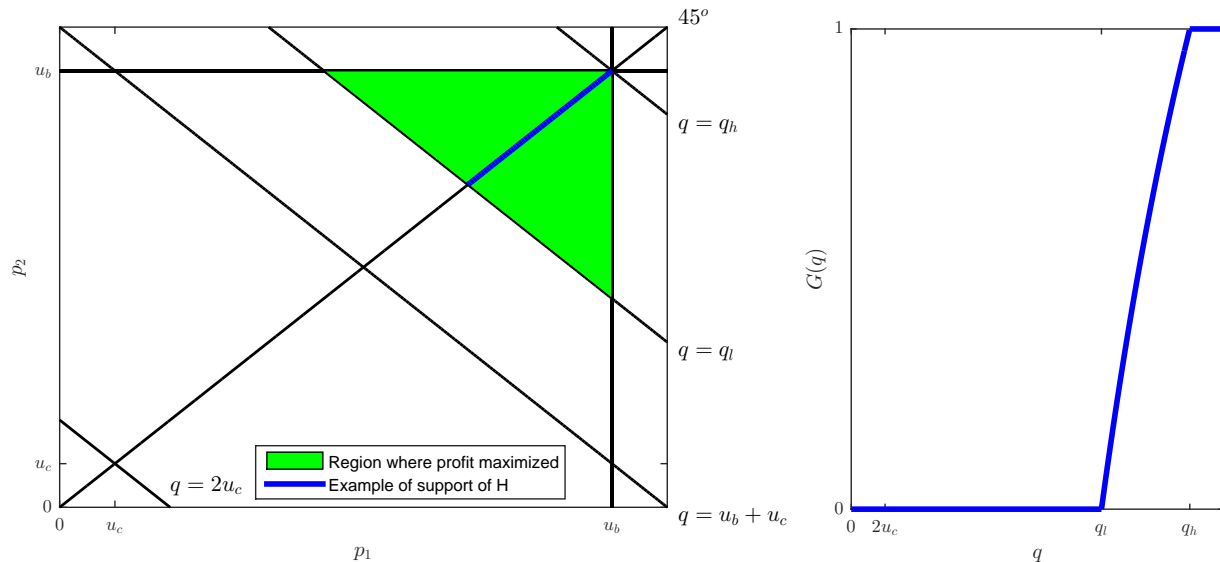
$$S(p_1, p_2) = \mu_b [\alpha + 2(1 - \alpha)(1 - G(p_1 + p_2))] (p_1 + p_2). \quad (11)$$

The seller is in the network of $\mu_b \alpha$ captive buyers of type b . A captive buyer of type b purchases both goods from the seller with probability 1, since both p_1 and p_2 are smaller than u_b . The seller is also in the network of $\mu_b 2(1 - \alpha)$ non-captive buyers of type b . A non-captive buyer of type b purchases both goods from the seller with probability $1 - G(p_1 + p_2)$, which is the probability that the second seller in the buyer's network has a basket price greater than $p_1 + p_2$. Finally, the seller is in the network of some buyers of type c . However, a buyer of type c never purchases from the seller, since both p_1 and p_2 are strictly greater than u_c . Notice that the seller's profit in (11) only depends on the sum of p_1 and p_2 . Therefore, the seller would attain the same profit by posting any other pair of prices $(p'_1, p'_2) \in R_1$ such that $p'_1 + p'_2 = p_1 + p_2$. ■

Lemma 4 is intuitive. Buyers of type b must make all of their purchases in the same location. Hence, they are indifferent between visiting a seller with prices $(p_1, p_2) \in R_1$ or a seller with prices $(p'_1, p'_2) \in R_1$ as long as the two sellers charge the same price for the whole basket of goods, i.e. $p_1 + p_2 = p'_1 + p'_2$. Buyers of type c never purchase from a seller with prices (p_1, p_2) nor from a seller with prices (p'_1, p'_2) , as both sellers charge prices that are above their willingness to pay. Since both buyers of type b and buyers of type c are indifferent between (p_1, p_2) and (p'_1, p'_2) , so is a seller. In turn, this implies that the profit of a seller with prices $(p_1, p_2) \in R_1$ can be written as a function of $q = p_1 + p_2$, i.e.,

$$S_1(q) = \mu_b [\alpha + 2(1 - \alpha)(1 - G(q))] q. \quad (12)$$

Now we are in the position to establish two properties of the distribution G of basket prices in a Bundled Equilibrium. First, the highest basket price, q_h , on the support of G equals $2u_b$. To see why, suppose that q_h is strictly smaller than $2u_b$. In this case, the profit for a seller with a basket price of q_h is then equal to $\mu_b \alpha q_h$, as this seller is the one with



Notes: This figure shows the possible range of the support of the joint distribution $H(p_1, p_2)$ and the shape of the cumulative distributions $G(q)$ in the Bundled Equilibrium.

Figure 3: Bundled Equilibrium

the highest basket price in the economy and, hence, only sells to captive buyers of type b . However, if the seller sets a basket price of $2u_b$, he attains a profit of $\mu_b \alpha 2u_b$, as the seller still only sells to captive buyers of type b . Since $\mu_b \alpha q_h < \mu_b \alpha 2u_b$, it follows that the seller with a basket price of q_h is not maximizing his profit, and, hence, this cannot be an equilibrium. Next, suppose that q_h is strictly greater than $2u_b$. In this case, Lemma 1 is violated. Hence, this cannot be an equilibrium either. Overall, in a Bundled Equilibrium, $q_h = 2u_b$.

Second, the support of G is an interval $[q_\ell, q_h]$. To see why, suppose that the support of G has a gap between the basket price q_0 and the basket price q_1 . In this case, a seller with a basket price of q_0 attains a profit of $\mu_b [\alpha + 2(1 - \alpha)(1 - G(q_0))] q_0$. A seller with a basket price of q_1 attains a profit of $\mu_b [\alpha + 2(1 - \alpha)(1 - G(q_1))] q_1$. Since G has a gap between q_0 and q_1 , $G(q_0) = G(q_1)$ and the seller with a basket price of q_0 makes the same number of trades as a seller with a basket price of q_1 but enjoys a lower profit per trade. Therefore, the seller with a basket price of q_0 does not maximize his profit, and, hence, this cannot be an equilibrium.

Next, we can solve for the distribution G of basket prices. In a Bundle Equilibrium, the seller's profit must attain its maximum for any q on the support of G . That is, $S_1(q) = S^*$ for all $q \in [q_\ell, q_h]$. Since $S_1(q_h) = S^*$ and $q_h = 2u_b$, the maximized profit S^* is given by $\mu_b \alpha 2u_b$. Since $S_1(q) = S^*$ for all $q \in [q_\ell, q_h]$ and $S^* = \mu_b \alpha 2u_b$, it follows that

$$\mu_b [\alpha + 2(1 - \alpha)(1 - G(q))] q = \mu_b \alpha 2u_b. \quad (13)$$

Solving (13) with respect to $G(q)$, we obtain the equilibrium distribution of basket prices

$$G(q) = 1 - \frac{\alpha}{2(1-\alpha)} \frac{2u_b - q}{q}. \quad (14)$$

Solving $G(q_\ell) = 0$ with respect to q_ℓ , we obtain the lower bound of the equilibrium distribution of basket prices

$$q_\ell = \frac{\alpha}{2-\alpha} 2u_b. \quad (15)$$

Figure 3 illustrates the equilibrium distribution of basket prices (14). Notice that the distribution of basket prices in (14) is the exactly the same as the equilibrium price distribution of Burdett and Judd (1983). This result is not surprising. In a Bundled Equilibrium, a seller and a buyer either trade the entire basket of goods or they do not trade at all. Hence, in a Bundled Equilibrium, our model of multiproduct retailing boils down to the single-product model of retailing of Burdett and Judd (1983), where the single product being traded in the market is the entire basket of goods.

In a Bundled Equilibrium, the distribution G of basket prices across stores is uniquely pinned down. The distribution of basket prices is non-degenerate, as there is a positive measure of sellers with any basket price q in the non-empty interval between q_ℓ and q_h . The fact that the distribution of basket prices is non-degenerate means that the equilibrium features price dispersion across stores, in the sense that some sellers are, on average, expensive and other sellers are, on average, cheap.

In a Bundled Equilibrium, the distribution H of price vectors across stores is not uniquely pinned down. In particular, any distribution of price vectors $H(p_1, p_2)$ that has support inside R_1 and that generates the distribution of basket prices $G(q)$ in (14) is such that every price on the support of H maximizes the profit of the seller and, hence, is an equilibrium. For instance, there is an equilibrium in which there are $G'(q)$ sellers with a basket price of q , and each one of them posts the price $q/2$ for both good 1 and good 2. This equilibrium does not feature relative price dispersion, as there is no dispersion across sellers in the price of a particular good at a particular seller relative to the average price charged by that seller. However, there always also exists an equilibrium in which there are $G'(q)$ sellers with a basket price of q , and each one of them posts prices $(p_1, q - p_1)$ where p_1 is randomly drawn from a uniform distribution with support $(u_c, q - u_c]$. In this equilibrium, there is relative price dispersion. Yet, such relative price dispersion is a matter of indifference, as neither buyers nor sellers care about the price of any individual good. In this sense, a Bundled Equilibrium—and, more generally, any model of multiproduct retailing in which buyers always purchase the entire basket of goods in the same location—does not offer a particularly compelling theory of relative price dispersion.

We conclude the analysis by identifying necessary and sufficient conditions for the existence of a Bundled Equilibrium. We find that this type of equilibrium exists if and only if

$$\frac{\mu_c}{\mu_b} \leq \frac{3\alpha - 2}{(2 - \alpha)u_c/u_b} - 1. \quad (16)$$

Condition (16) is satisfied if: (i) the market is not too competitive, in the sense that the fraction α of buyers who are in contact with only one seller is greater than $2/3$; and (ii) the relative number of type- c buyers, μ_c/μ_b , and/or the relative willingness to pay of type- c buyers, u_c/u_b , is not too large. Intuitively, if the market is too competitive, some sellers would want to set a price q for the basket of goods that is smaller than $u_b + u_c$. Hence, a Bundled Equilibrium would not exist. If the relative number of type- c buyers and/or the relative willingness to pay of type- c buyers is too high, some sellers would want to trade with type- c buyers and set a price p for one or both goods that is smaller than u_c . Hence, a Bundled Equilibrium would not exist.

The following proposition summarizes the characterization of a Bundled Equilibrium.

Proposition 1: (Bundled Equilibrium) *(i) In a Bundled Equilibrium, the distribution of basket prices, G , is continuous over the support $[q_\ell, q_h]$ and given by (14); the distribution of price vectors across sellers, H , is not uniquely pinned down. (ii) A Bundled Equilibrium exists if and only if (16) is satisfied.*

Proof: Appendix C. ■

2.4 Discrimination Equilibrium

In this section, we consider an equilibrium in which some sellers have a basket price q greater than $u_b + u_c$, and some sellers have a basket price smaller than $u_b + u_c$ and greater than $2u_c$. We refer to this type of equilibrium as a *Discrimination Equilibrium*, as in this equilibrium some sellers set their prices so as to discriminate between the high-valuation buyers who must purchase all the goods in the same location and the low-valuation buyers who can purchase different goods in different locations. In the first part of the section, we characterize the price distribution in a Discrimination Equilibrium and show that it necessarily features relative price dispersion. In the second part, we identify necessary and sufficient conditions for the existence of this type of equilibrium.

We start the characterization of a Discrimination Equilibrium by focusing on the sellers with a basket price q strictly greater than $u_b + u_c$. Following the same arguments as in the previous section, we can easily show that, among sellers with $q > u_b + u_c$, the distribution of basket prices $G(q)$ has support over the interval $[q^*, q_h]$, with $u_b + u_c < q^* < q_h = 2u_b$.

Moreover, using the fact that all of the sellers with $q \in [q^*, q_h]$ must attain the maximized profit S^* , we can easily show that the maximized profit S^* is equal to $\mu_b \alpha 2u_b$ and that, for all $q \in [q^*, q_h]$, the distribution of basket prices $G(q)$ is equal to

$$G(q) = 1 - \frac{\alpha}{2(1-\alpha)} \frac{2u_b - q}{q}. \quad (17)$$

Next, we focus on sellers with a basket price q smaller than $u_b + u_c$ and strictly greater than $2u_c$. In Lemma 2, we established that a seller with a basket price $q \in (2u_c, u_b + u_c]$ sets the price of one good below the willingness to pay of type- c buyers, and sets the price of the other good in between the willingness to pay of type- c buyers and the willingness to pay of type- b buyers. That is, a seller with a basket price $q \in (2u_c, u_b + u_c]$ sets prices $(p_1, p_2) \in R_2$. Now, consider a seller who posts prices $(p_1, p_2) \in R_2$. The seller trades the basket of goods to buyers of type c at the price $q = p_1 + p_2$, and he trades the cheaper of the two goods (say good i) to buyers of type b at the price p_i . Thus, the seller attains a profit of

$$\begin{aligned} S_{2i}(q, p_i) &= \mu_b [\alpha + 2(1-\alpha)(1-G(q))] q \\ &\quad + \mu_c [\alpha + 2(1-\alpha)(1-F_i(p_i))] p_i. \end{aligned} \quad (18)$$

Note that (18) makes use of the fact that $G(q)$ does not have any mass points, and $F_i(p)$ does not have any mass points over the interval $(0, u_c]$.

The next lemma shows that, for all $p \in [0, u_c]$, the fraction of sellers charging less than p for good 1 is exactly the same as the fraction of sellers charging less than p for good 2. That is, $F_1(p) = F_2(p) = F(p)$ for all $p \in [0, u_c]$. The lemma implies that the profit of a seller in region R_2 is symmetric in the two goods. That is, $S_{21}(q, p) = S_{22}(q, p) = S_2(q, p)$. The lemma is intuitive. If $F_1(p) > F_2(p)$ for $p \in (p_0, p_1)$, with $0 \leq p_0 < p_1 \leq u_c$, then a seller posting the prices $(p, q-p) \in R_2$ would be better off posting the prices $(q-p, p)$ instead. In fact, the seller trades the basket of goods to the same number of type- b buyers and at the same price by posting either $(q-p, p)$ or $(p, q-p)$. However, by posting $(q-p, p)$ rather than $(p, q-p)$, the seller trades the cheaper good to more type- c buyers even though he charges the same price for it. Hence, if $F_1(p) > F_2(p)$ for $p \in (p_0, p_1)$, all sellers posting the prices $(p, q-p) \in R_2$ would be better off switching the price tags of the two goods until $F_1(p) = F_2(p)$.

Lemma 5: *In a Discrimination Equilibrium, $F_1(p) = F_2(p)$ for all $p \in [0, u_c]$.*

Proof: Appendix C. ■

The next lemma shows that the profit of a seller in region R_2 attains the maximum S^* for all prices of the basket q and prices of the cheaper good p such that q is in the interval $[q_\ell, u_b + u_c]$ and p is in the interval $[p_\ell, u_c]$, where q_ℓ denotes the lower bound on the

support of the price distribution of baskets and p_ℓ denotes the lower bound on the support of the price distribution of an individual good. That is, $S_2(q, p) = S^*$ for all (q, p) such that $q \in [q_\ell, u_b + u_c]$ and $p \in [p_\ell, u_c]$. The proof of the lemma follows the same strategy as Proposition 3 in Menzio and Trachter (2015b). The gist of the proof is to show that, if profits are not constant for all (q, p) such that $q \in [q_\ell, u_b + u_c]$ and $p \in [p_\ell, u_c]$, there are either gaps on the support of the distribution of G over the interval $[q_\ell, u_b + u_c]$ or gaps on the support of the distribution F over the interval $[p_\ell, u_c]$. In turn, if there are gaps on the support of one of the two distributions, there are some sellers that could increase their profits by either increasing the price of the basket or by increasing the price of one of the cheaper good.

Lemma 6: *In a Discrimination Equilibrium, $S_2(q, p) = S^*$ for all (q, p) such that $q \in [q_\ell, u_b + u_c]$ and $p \in [p_\ell, u_c]$.*

Proof: Appendix C. ■

We are now in the position to solve for the lowest basket price q^* posted by sellers in region R_1 , for the marginal distribution $G(q)$ of basket prices among sellers in region R_2 , and for the marginal distribution $F(p)$ of prices among sellers in region R_2 . Lemma 6 implies that a seller posting prices (u_c, u_b) attains the maximized profit S^* , i.e.,

$$\mu_b [\alpha + 2(1 - \alpha) (1 - G(u_b + u_c))] (u_b + u_c) + \mu_c [\alpha + 2(1 - \alpha) (1 - F(u_c))] u_c = S^*. \quad (19)$$

Similarly, a seller posting prices (p_1, p_2) such that $p_1 \in (u_c, u_b]$, $p_2 \in (u_c, u_b]$, and $p_1 + p_2 = q^*$ attains the maximized profit S^* , i.e.,

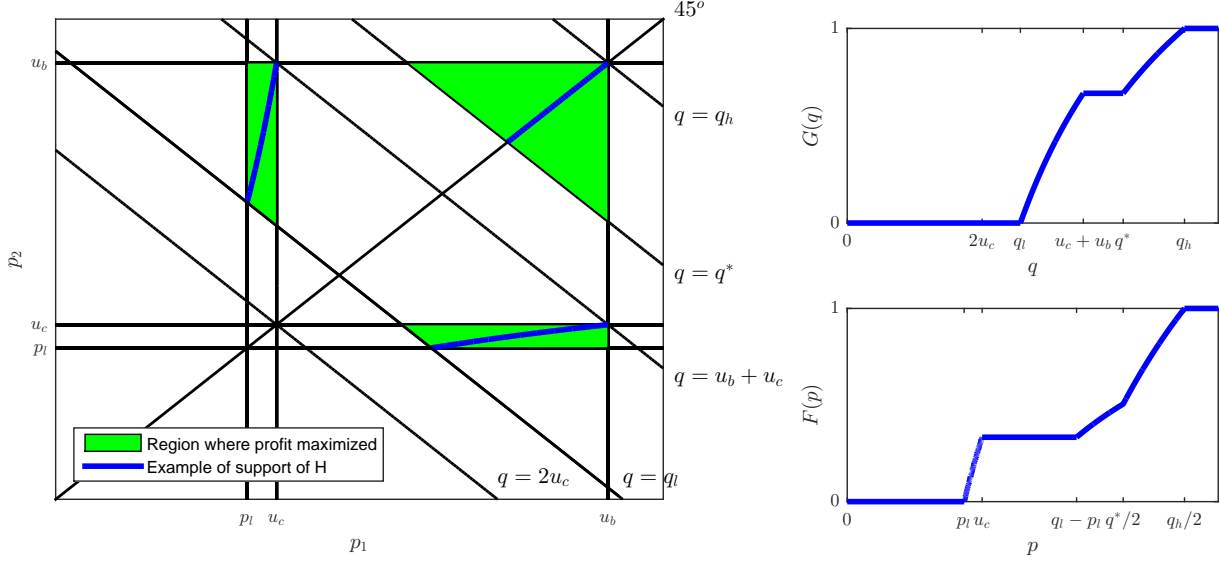
$$\mu_b [\alpha + 2(1 - \alpha) (1 - G(q^*))] q^* = S^*. \quad (20)$$

Notice that the fraction of sellers with a basket price smaller than q^* is the same as the fraction of sellers with a basket price smaller than $u_b + u_c$, i.e., $G(q^*) = G(u_b + u_c)$. Also, notice that the fraction of sellers that charge less than u_c for good 1 is half of the fraction of sellers with a basket price smaller than q^* , i.e., $F(u_c) = G(q^*)/2$. Using these two observations and the fact that the left-hand side of (19) is equal to the left-hand side of (20), we obtain

$$\begin{aligned} & \mu_b [\alpha + 2(1 - \alpha) (1 - G(q^*))] (u_b + u_c) + \mu_c [\alpha + 2(1 - \alpha) (1 - G(q^*)/2)] u_c \\ & = \mu_b \{ \alpha + 2(1 - \alpha) [1 - G(q^*)] \} q^*. \end{aligned} \quad (21)$$

Equation (21) can be solved with respect to q^* to obtain

$$q^* = \frac{2\alpha(1 + u_c/u_b) + \alpha(\mu_c/\mu_b)(u_c/u_b)}{4\alpha - (2 - \alpha)(\mu_c/\mu_b)(u_c/u_b)} 2u_b. \quad (22)$$



Notes: This figure shows the possible range of the support of the joint distribution $H(p_1, p_2)$, the shape of the cumulative distributions $G(q)$, and an example of the shape of the cumulative distribution $F(p)$ in the Discrimination Equilibrium.

Figure 4: Discrimination Equilibrium

Lemma 6 implies that a seller posting any prices (p_1, p_2) such that $p_2 \in (u_c, u_b]$ and $q = p_1 + p_2 \in [q_\ell, u_b + u_c]$ attains the same profit as a seller posting prices (u_c, u_b) , i.e.,

$$\begin{aligned} & \mu_b [\alpha + 2(1 - \alpha) (1 - G(q))] q + \mu_c [\alpha + 2(1 - \alpha) (1 - F(u_c))] u_c \\ &= \mu_b [\alpha + 2(1 - \alpha) (1 - G(u_b + u_c))] (u_b + u_c) + \mu_c [\alpha + 2(1 - \alpha) (1 - F(u_c))] u_c. \end{aligned} \quad (23)$$

Using the fact that $G(u_b + u_c) = G(q^*)$ and solving (23) with respect to $G(q)$, we find that the distribution of basket prices for $q \in [q_\ell, u_b + u_c]$ is given by

$$G(q) = G(q^*) - \frac{\alpha + 2(1 - \alpha) (1 - G(q^*)) u_b + u_c - q}{2(1 - \alpha) q}. \quad (24)$$

Solving the equation $G(q_\ell) = 0$ with respect to q_ℓ , we find that the lowest price on the support of the distribution of basket prices is given by

$$q_\ell = \frac{2\alpha u_b u_b + u_c}{2 - \alpha q^*}. \quad (25)$$

Lemma 6 also implies that a seller posting prices (p_1, p_2) such that $p_1 \in [p_\ell, u_c]$, $p_2 \in (u_c, u_b]$ and $p_1 + p_2 = q_\ell$ attains the same profit as a seller posting prices $(u_c, q_\ell - u_c)$, i.e.

$$\begin{aligned} & \mu_b [\alpha + 2(1 - \alpha)] q_\ell + \mu_c [\alpha + 2(1 - \alpha) (1 - F(p))] p \\ &= \mu_b [\alpha + 2(1 - \alpha)] q_\ell + \mu_c [\alpha + 2(1 - \alpha) (1 - F(u_c))] u_c. \end{aligned} \quad (26)$$

Using the fact that $F(u_c) = G(q^*)/2$ and solving the equation (26) with respect to $F(p)$, we find that the distribution of good-1 prices for $p \in [p_\ell, u_c]$ is given by

$$F(p) = \frac{G(q^*)}{2} - \frac{\alpha + 2(1 - \alpha)(1 - G(q^*)/2)}{2(1 - \alpha)} \frac{u_c - p}{p}. \quad (27)$$

Solving the equation $F(p_\ell) = 0$ with respect to p_ℓ , we find that the lowest price on the support of the distribution of good-1 prices is given by

$$p_\ell = \frac{\alpha + 2(1 - \alpha)[1 - G(q^*)/2]}{2 - \alpha} u_c.$$

This completes the characterization of a Discrimination Equilibrium. In this type of equilibrium, there is a group of sellers who sets a basket price of $q \in [q^*, q_h]$ and the prices p_1 and p_2 in between u_c and u_b . These sellers trade (with some probability) the basket of goods to buyers of type b and never trade with buyers of type c . There is also a group of sellers who set a basket price of $q \in [q_\ell, u_b + u_c]$. Half of these sellers set p_1 below u_c and p_2 between u_c and u_b . These sellers trade (with some probability) the whole basket of goods to buyers of type b and good 1 to buyers of type c . The other half of the sellers sets p_2 below u_c and p_1 between u_c and u_b . These sellers trade (with some probability) the whole basket of goods to buyers of type b and good 2 to buyers of type c . There are no sellers who set a basket price of q in the interval $(u_b + u_c, q^*)$.

The distribution of basket prices $G(q)$ is given by (17) for $q \in [q^*, q_h]$ and by (24) for $q \in [q_\ell, u_b + u_c]$. The distribution $G(q)$ is such that the seller's profit from trading the basket of goods to buyers of type b is equal to S^* for all $q \in [q^*, q_h]$, and it is equal to $S^* - \mu_c[\alpha + 2(1 - \alpha)(1 - F(u_c))]u_c$ for all $q \in [q_\ell, u_b + u_c]$. The distribution $G(q)$ has a gap between $u_b + u_c$ and q^* . The gap exists because a seller with a basket price of $u_b + u_c$ trades with both buyers of type b and buyers of type c , while a seller with a basket price greater than $u_b + u_c$ only trades with buyers of type b . Therefore, a seller strictly prefers setting a basket price of $u_b + u_c$ rather than setting any basket price just above $u_b + u_c$. The distribution of prices for an individual good $F(p)$ is given by (27) for $p \in [p_\ell, u_c]$. The distribution $F(p)$ is such that the seller's profit from trading the cheaper good to buyers of type c is equal to $S^* - \mu_b(2 - \alpha)q_\ell$ for all $p \in [p_\ell, u_c]$. The distribution $F(p)$ is not uniquely pinned down for $p \in (u_c, u_b]$. Intuitively, this is the case because a seller who charges a price of $p > u_c$ for one good only trades that good to buyers of type b together with the other good.

The distribution of price vectors H is not uniquely pinned down. For sellers with a basket price $q \in [q^*, q_h]$, any distribution H that has support inside R_1 and that generates the marginal distribution of basket prices $G(q)$ in (17) is consistent with equilibrium. For

example, there is an equilibrium in which, for all $q \in [q^*, q_h]$, there are $G'(q)$ sellers with a basket price of q and each of them posts the prices $(q/2, q/2)$. For sellers with a basket price $q \in [q_\ell, u_b + u_c]$, any distribution H that has support inside R_2 and that generates the marginal distribution of basket prices $G(q)$ in (24) and the marginal distribution of individual good prices $F(p)$ in (27) is consistent with equilibrium. For example, there is an equilibrium in which, for all $p \in [p_\ell, u_c]$, $2F'(p)$ sellers have a basket price of $\phi(p)$, $F'(p)$ sellers post the prices $(p, \phi(p) - p)$, and $F'(p)$ sellers post the prices $(\phi(p) - p, p)$, where

$$\phi(p) = \frac{[\alpha + 2(1 - \alpha)(1 - G(q^*))](u_b + u_c)}{[\alpha + 2(1 - \alpha)(1 - G(q^*))] + 2[\alpha + 2(1 - \alpha)(1 - G(q^*)/2)](u_c - p)/p}. \quad (28)$$

A Discrimination Equilibrium features price dispersion across sellers, in the sense that some sellers are on average more expensive, while some sellers are on average cheaper. This property of equilibrium follows immediately from the fact that the distribution of basket prices is non-degenerate. A Discrimination Equilibrium always features relative price dispersion, in the sense that there is variation across sellers in the price of a particular good at a particular seller relative to the average price charged by that seller. This property of equilibrium follows immediately from the fact that half of the sellers with a basket price $q \in [q_\ell, u_b + u_c]$ have a relative price for good 1 that is strictly greater than 1, while the other half of the sellers with a basket price $q \in [q_\ell, u_b + u_c]$ have a relative price for good 1 that is strictly smaller than 1.

Let us briefly explain why relative price dispersion must emerge in equilibrium. Competition between sellers drives part of the distribution of basket prices to the region where q is between $2u_c$ and $u_b + u_c$. A seller with a basket price between $2u_c$ and $u_b + u_c$ never finds it optimal to post the same price for both goods. Instead, the seller finds it optimal to set the price of one good below and the price of the other good above the willingness to pay of type- c buyers. That is, a seller with a basket price q between $2u_c$ and $u_b + u_c$ finds it optimal to follow an asymmetric pricing strategy for the two goods. However, if some sellers post a higher price for good 1 than for good 2, other sellers must post a higher price for good 2 than for good 1, or else there would be some unexploited profit opportunities. That is, the distribution of prices for the two goods must be symmetric across sellers with a basket price q between $2u_c$ and $u_b + u_c$. The asymmetric pricing strategy followed by each individual seller combined with the symmetry of the price distribution across sellers implies relative price dispersion.

Sellers follow an asymmetric pricing strategy to discriminate between the two types of buyers. The difference in the willingness to pay of type- b and type- c buyers gives sellers a desire to price discriminate. The difference in the ability of type- b buyers and type- c

buyers to purchase different items in different locations gives sellers the opportunity to price discriminate. In fact, by pricing the two goods asymmetrically, a seller can charge a high average price to the high-valuation buyers who need to purchase all the items together (the buyers of type b) and charge a low price for one good to the low-valuation buyers who can purchase different items at different locations (the buyers of type c).

It is interesting to contrast the type of price discrimination described above with intertemporal price discrimination (see, e.g., Conlisk, Gerstner and Sobel 1984 and Sobel 1984 or, in a search-theoretic context, Albrecht, Postel-Vinay and Vroman 2013 and Menzio and Trachter (2015b)). The key to intertemporal price discrimination is a negative correlation between a buyer's valuation and his ability to intertemporally substitute purchases. A seller can exploit this negative correlation by having occasional sales. The low-valuation buyers, who are better able to substitute purchases intertemporally, will take advantage of the sales and will end up paying low prices. The high-valuation buyers, who are unable to substitute purchases intertemporally, will not take advantage of the sales and will end up paying high prices. In contrast, our theory of price discrimination is based on a negative correlation between a buyer's valuation and his ability to shop in multiple stores. Moreover, while intertemporal price discrimination takes the form of time variation in the price of the same good, our theory of price discrimination takes the form of variation in the price of different goods relative to the average store price.

We conclude the analysis by identifying necessary and sufficient conditions for the existence of a Discrimination Equilibrium. We find that this type of equilibrium exists if and only if

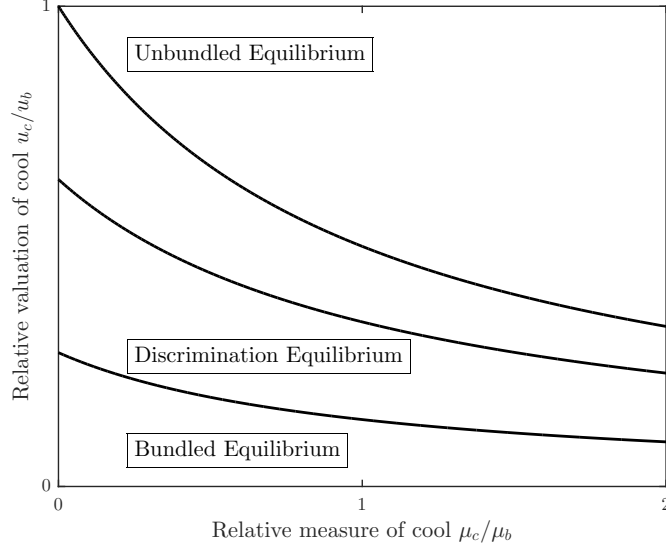
$$\frac{\mu_c}{\mu_b} > \frac{3\alpha - 2}{(2 - \alpha)u_c/u_b} - 1 \quad (29)$$

and

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha - (2 - \alpha)u_c/u_b}{1 + (2 - \alpha)u_c/u_b} \frac{1 + u_c/u_b}{u_c/u_b}. \quad (30)$$

Condition (29) guarantees that some sellers find it optimal to post basket prices below $u_b + u_c$. The condition is satisfied if: (i) the market is sufficiently competitive, in the sense that the fraction α of buyers who are in contact with only one seller is smaller than $2/3$; or (ii) the relative number of type- c buyers, μ_c/μ_b , and/or the relative willingness to pay of type- c buyers, u_c/u_b , is large enough. Condition (30) guarantees that no seller finds it optimal to post prices below $2u_c$. The condition is satisfied if: (i) the market is not too competitive, in the sense that the fraction α of buyers who are in contact with only one seller is greater than $2(u_c/u_b)/(1 + 2(u_c/u_b))$; (ii) the relative number of type- c buyers, μ_c/μ_b , and/or the relative willingness to pay of type- c buyers, u_c/u_b , are low enough. As we shall see in the next section, conditions (29) and (30) define a non-empty parameter space.

Figure 5: Equilibrium type depending on buyers in market



Notes: This figure illustrates how the type of equilibrium depends on the shares and relative valuations of the two types of buyers in the market.

The following proposition summarizes the characterization of a Discrimination Equilibrium.

Proposition 2: (Discrimination Equilibrium) (i) In a Discrimination Equilibrium, the distribution of basket prices, G , is continuous over the support $[q_\ell, u_b + u_c] \cup [q^*, q_h]$, and it is given by (17) for $q \in [q^*, q_h]$ and by (24) for $q \in [q_\ell, u_b + u_c]$; the distribution of prices for good i , F_i , is continuous over the interval $[p_\ell, u_c]$ and it is given by (27) for both $i = 1$ and 2 ; the distribution of price vectors, H , is not uniquely pinned down. (ii) A Discrimination Equilibrium exists if and only if (29) and (30) are satisfied.

Proof: Appendix C. ■

2.5 Other Equilibria

A Bundled Equilibrium exists if and only if the relative measure of type- c buyers, μ_c/μ_b , is smaller than the right-hand side of (16). The lowest curve in Figure 5 is the plot of the right-hand side of (16) in the $\{u_c/u_b, \mu_c/\mu_b\}$ space. Hence, a Bundle Equilibrium exists if and only if the parameter values lie below the bottom curve.¹¹ In this type of equilibrium, every seller sets a basket price q strictly greater than $u_b + u_c$ and posts prices for the individual goods that are strictly greater than the willingness to pay of type- c buyers and smaller than

¹¹It is easy to verify that, when condition (16) is satisfied, the only equilibrium is a Bundle Equilibrium.

the willingness to pay of type- b buyers. Every seller trades the entire basket of goods to buyers of type b and trades nothing to buyers of type c . The equilibrium distribution of basket prices is the same as in the single-product model of retailing by Burdett and Judd (1983).

A Discrimination Equilibrium exists if and only if the relative measure of type- c buyers, μ_c/μ_b , is greater than the right-hand side of (29) and smaller than the right-hand side of (30). Notice that the right-hand side of (29) is equal to the right-hand side of (16), which is the lowest curve in Figure 5. Also, notice that the right-hand side of (29) is strictly smaller than the right-hand side of (30), which is the middle curve in Figure 5. Hence, a Discrimination Equilibrium exists if and only if the parameter values fall in the non-empty region between the bottom and the middle curves.¹² In this type of equilibrium, there are some sellers with a basket price $q > u_b + u_c$ who post prices p_1 and $p_2 \in (u_c, u_b]$. These sellers trade the basket of goods to buyers of type b and trade nothing to buyers of type c . There are also some sellers with a basket price $q \in (2u_c, u_c + u_b]$ who set the price of one good above and the price of the other good below u_c . These sellers trade the basket of goods to buyers of type b and the cheaper good to buyers of type c .

If the relative measure of type- c buyers is greater than the right-hand side of (30), other more complex types of equilibria emerge. In this paper, we do not wish to analyze these equilibria in detail, but we still find it instructive to describe some of their properties. Any equilibrium is an *Unbundled Equilibrium* when and only when

$$\frac{\mu_c}{\mu_b} > \frac{1 - u_c/u_b}{u_c/u_b}. \quad (31)$$

Notice that the right-hand side of (30) is strictly smaller than the right-hand side of (31), which is the top curve in Figure 5. Hence, an equilibrium is unbundled if and only if the parameter values fall in the region above the top curve. In this type of equilibrium, every seller sets a basket price q smaller than $2u_c$ and posts prices for the individual goods that are smaller than u_c . Every seller trades the basket of goods to buyers of type b and either the basket of goods or one of the two individual goods to buyers of type c .

If the relative measure of type- c buyers lies between the top and the middle curves in Figure 5, the equilibrium is a combination between a Discrimination and an Unbundled Equilibrium. In this type of equilibrium, there is a first group of sellers with a basket price $q > u_b + u_c$, a second group of sellers with a basket price $q \in (2u_c, u_b + u_c]$, and a third group of sellers with a basket price $q \leq 2u_c$. Sellers in the first group trade the basket of goods to buyers of type b and nothing to buyers of type c . Sellers in the second group trade the

¹²When conditions (29) and (30) are satisfied, the only equilibrium is a Discrimination Equilibrium.

basket of goods to buyers of type b and the cheaper good to buyers of type c . Sellers in the third group trade the basket of goods to buyers of type b and either both or either one of the individual goods to buyers of type c .

3 Calibration and Validation

In this section, we calibrate and validate our theory of relative price dispersion. In Section 3.1, we calibrate the theory using data on the extent and sources of dispersion in the price at which the same good is sold by different sellers in the same market and week, the number of stores visited by different households, the relationship between the prices paid by a household for the same basket of goods in the same market and quarter, and the number of stores from which the households shops in a given quarter. In Section 3.2, we compare the predictions of the calibrated theory on the extent and sources of dispersion in the prices paid by different households for the same basket of goods in the same market and quarter. We find that our theory of relative price dispersion is consistent not only with the key facts about price dispersion but also with the main differences between the extent and sources of price dispersion across stores and price dispersion across households.

3.1 Calibration

For the purposes of our quantitative analysis, we need to consider a dynamic version of the static model of Section 2. We assume that the type of the buyer remains constant over time. However, a buyer may change his network of sellers from one period to the next because, for example, he moves from one part of the city to another, he changes his job, and so on. In particular, we assume that a buyer keeps the same network of sellers from one period to the next with probability ρ , and he samples a new network of sellers with probability $1 - \rho$. Conditional on changing his network, the buyer contacts one randomly selected seller with probability α and two randomly selected sellers with probability $1 - \alpha$. Sellers post prices in every period. Since the probability that a buyer changes network is independent of his current network, the pricing problem of the sellers in every period is the same as in the static model of Section 2. Hence, in the presence of any positive menu cost, sellers find it strictly optimal to keep their prices constant from one period to the next.

When quantifying the model, we have to deal with the fact that the theory does not uniquely pin down the distribution H of price vectors across sellers. As explained in Section 2, the theory uniquely pins down the distribution G of basket prices and the distribution F of prices of individual goods, but it does not uniquely pin down H . Yet, we find it

natural to assume that: (i) sellers with a basket price of $q \in (u_b + u_c, 2u_b]$ set the same price for both goods; and (ii) sellers with a basket price of $q \in (2u_c, u_b + u_c]$ set the prices $(\phi^{-1}(q), q - \phi^{-1}(q))$ with probability 1/2 and the prices $(q - \phi^{-1}(q), \phi^{-1}(q))$ with probability 1/2. Under conditions (i) and (ii), the equilibrium is *symmetric*, in the sense that $H(p_1, p_2) = H(p_2, p_1)$, and *rank-preserving*, in the sense that the rank of a seller in the distribution of basket prices is the same as the rank of a seller in the distribution of the lowest prices of an individual good.¹³ There are other equilibria that are symmetric and rank-preserving. However, conditions (i) and (ii) select the one with the lowest relative price dispersion.¹⁴

The dynamic version of the model has six parameters: the measure μ_b of buyers of type b , the measure μ_c of buyers of type c , the valuation u_b of buyers of type b , the valuation u_c of buyers of type c , the fraction α of buyers that have only one seller in their network, and the probability ρ that buyers keep the same network of sellers from one week to the next.

The equilibrium of the model depends on the ratio between the measure of type- c buyers and the measure of type- b buyers, but not on the two measures separately. Hence, we can normalize μ_b to 1. Similarly, the equilibrium of the model depends on the ratio between the valuation of type- c buyers and the valuation of type- b buyers, but not on the two valuations separately. Hence, we can normalize u_b to 1. We calibrate the remaining four parameters so as to match four moments of the data. First, we target a measure of the dispersion of prices for the same good in the same market and in the same week. Second, we target a measure of the fraction of price dispersion that is due to differences in the store component of the price and the fraction of price dispersion that is due to differences in the store-good component of the price. Third, we target a measure of the effect of shopping from an additional store on the prices paid by a household. Finally, we target a measure of the number of different stores visited by a given household.

As discussed in Section 2, we interpret the model as a theory of the persistent component of prices at different stores. For this reason, we want the model to match the variance of prices that is caused by persistent price differences across stores. As documented in Section 1, the standard deviation of prices for the same good in the same market and in the same

¹³It is immediate to see that the equilibrium is symmetric under (i) and (ii). To see that the equilibrium is rank-preserving, notice that, for sellers with $q \in [q_\ell, u_b + u_c]$, the function $\phi^{-1}(q)$ relating the price of the seller's basket to the price of the seller's cheapest good is such that $F(\phi^{-1}(q)) = G(q)/2$ and, hence, the rank of the seller in the distribution of basket prices is $G(q)$ and the rank of the seller in the distribution of the lowest price for an individual good is $2F(\phi^{-1}(q)) = G(q)$. For sellers with $q \in [q^*, 2u_b]$, the rank of the seller in the distribution of basket prices is $G(q)$ and the rank of the seller in the distribution of the lowest price for an individual good is $2F(u_c) + G(q) - G(q^*) = G(q)$.

¹⁴The other symmetric, rank-preserving equilibria are such that half of the sellers with a basket price of $q \in [q^*, 2u_b]$ post prices $(\psi(q), q - \psi(q))$ and the other half post prices $(q - \psi(q), \psi(q))$, with $\psi(q) \in (u_c, q/2]$. Clearly, conditions (i) and (ii) select the symmetric, rank-preserving equilibrium with the lowest amount of relative price dispersion.

week is 15.3%, and the variance is 2.34%. The fraction of the variance of prices that is due to differences in the store component of prices is 15.5%. Nearly all of the differences in the store component of prices are persistent, in the sense that they are either due to differences in the fixed effect or to differences in the AR part of the store component. The fraction of the variance of prices that is due to differences in the store-good component of prices is 84.5%. Approximately 36% of the differences in the store-good component of prices are persistent, and 64% are transitory. Based on these observations, the variance of prices that comes from persistent differences in prices is 1% (and the standard deviation is 10%), the fraction of this variance due to persistent differences in the store component is 34% percent, and the variance due to persistent differences in the store-good component is 66% percent.

In order to obtain a measure of the effect of shopping at more stores on the prices paid by households, we use the Kilts Nielsen Consumer Panel (KNCP), which tracks the shopping behavior of approximately 50,000 households over the period from 2004 to 2009. Households are drawn from 54 geographically dispersed markets, known as Scantrack markets, each of which roughly corresponds to a metropolitan statistical area. Demographic data on panelists are collected at the time of entry into the panel and updated annually through a written survey. Panelists provide information about each of their shopping trips using a Universal Product Code (UPC) scanning device. More specifically, when a panelist returns from a shopping trip, he uses the device to enter details about the trip, including the date and the store where the purchases were made. The panelist then scans the barcode of the purchased good and enters the number of units purchased. The price of the good is recorded either automatically or manually depending on whether the store where the good was purchased is covered by Nielsen or not.

We follow the methodology in Aguiar and Hurst (2007) to compute a price index for every household in the KNCP. Specifically, we define the price index of household i in market m and quarter t as the ratio between the dollar amount that the household paid to purchase its basket of goods and the amount that the household would have spent had it paid, for each good in its basket, the average price of that good in market m and quarter t . We refer the reader to Kaplan and Menzio (2015) for further details about the construction of price indexes and for the sample selection criteria.

In Panel (a) of Table 3, we report the results of a regression of the price index of household i in quarter t on the log (or the level) of the ratio between the number of different stores in which household i shopped during quarter t and the dollar expenditures of household i in quarter t . In all specifications, we control for household size, the age and education of household members, and we add market dummies. In some specifications, we control for

Table 3: Regression of household price indexes on indicators of multistop shopping


(a)	Log	Log	Level	Level
Stores/Expenditure	-0.01124** (0.00041)	-0.01291** (0.00041)	-0.09099** (0.00363)	-0.10365** (0.00349)
FE	No	Yes	No	Yes
R^2	0.01776	0.0074	0.02041	0.01523
(b)	Log	Log	Level	Level
Stores	-0.03424** (0.00072)	-0.01577** (0.00051)	-0.01196** (0.00027)	-0.00367** (0.00018)
Expenditure	0.00964** (0.00041)	0.01261** (0.00044)	0.00008** (0.00000)	0.00007** (0.00000)
FE	No	Yes	No	Yes
R^2	0.02875	0.00753	0.02707	0.00235

Notes: This table presents results for regressions of household price indexes on indicators of multistop shopping: the average number of different stores visited per dollar spent in the quarter (Panel A) and number of different stores visited per quarter, conditioning on dollars spent per quarter (Panel B). The level models have expenditures in levels and log models have expenditures in logs. In all regressions: N=880104, clusters=78758.

household fixed effects, while in other specifications we do not. Column 1 reports the results of the regression on the log of stores-per-dollar without household fixed effects. It shows that a household that visits twice as many stores per dollar spent has a price index that is 1.12% lower. Column 2 reports the results of the regression on the log of stores-per-dollar with household fixed effects. It shows that, in quarters when a household visits twice as many stores per dollar spent, it enjoys a 1.29% lower price index. Columns 3 and 4 report the results of the regressions on the level of stores-per-dollar when household fixed effects are respectively excluded and included. For all specifications, the regression coefficient on stores per dollar is negative and significant. In Panel (b), we carry out the same regression using the log (or the level) of the number of different stores in which the households shopped during quarter t . We find that the regression coefficient on the number of stores is always negative and significant. Since our model has only limited heterogeneity across buyers, we choose as a calibration target the regression coefficient on the log of stores-per-dollar.

In Figure 6, we display the distribution of the number of stores visited in a quarter by households in KNCP. Approximately 32% of households do all of their shopping in a quarter at one store. Approximately 28% of households do all of their shopping in a quarter at two stores. Only 40% of households shop from more than two stores in a quarter. As a

Figure 6: Number of different stores visited by households



Notes: This figure displays a histogram of the number of different stores visited by households within a quarter.

calibration target for the model, we choose the average number of stores, 2.48, from which a household shops in a quarter.

We calibrate the parameters of the model by minimizing the sum of the absolute value of the percentage deviation between each of the targeted moments and its counterpart in the model. Table 4 reports the targeted moments, the value of these moments in the calibrated model, and the calibrated value of the parameters.

3.2 Validation

Our theory of shopping and pricing in the retail market—albeit rather stylized—matches the extent and sources of dispersion in the price of the same good across different sellers. Indeed, the calibrated model matches the standard deviation of prices posted by different sellers for the same good in the same market and in the same week. Also, the calibrated model matches quite well the fraction of the variance of prices posted by different sellers that is due to differences in the store component of the price (37% in the model, 34% in the data) and the fraction that is due to differences in the store-good component of the price (63% in the model, 66% in the data). The theory’s explanation for this decomposition is simple. There is dispersion in the store component of prices for the same reasons as in Burdett and Judd (1983). Specifically, as some buyers have only one seller in their network and other buyers have multiple sellers in their network, sellers must post different basket prices in equilibrium. There is dispersion in the store-good component of prices (i.e., relative

Table 4: Calibration

Targets	Data	Model
Standard deviation of prices	10%	10%
Share of variance of prices due to store component	34%	37%
Average number of stores visited	2.48	2.41
Regression coefficient of price index on store number	-1.3%	-1.9%
Additional moments		
Standard deviation of price index	7.8%	6.6%
Share of variance of price index due to store component	55%	71%

Parameter	μ_c/μ_b	u_c/u_b	α	ρ
Value	0.041	0.68	0.89	0.87

price dispersion) because of price discrimination. Specifically, as high-valuation buyers need to purchase all the goods from the same retailer and low-valuation buyers can purchase different goods from different retailers, an individual seller wants to set asymmetric prices for the two goods. And, in equilibrium, for every seller posting a higher price for the first good than for the second, there must be another seller doing the opposite.

Our theory matches the negative relationship between the price index of a household and the number of stores from which the household shops during a quarter (normalized by the household's expenditures). This is intuitive. According to the model, households who shop from more sellers are more likely to be buyers who are willing to go through the stores in their network to purchase different goods at the lowest price. Households who shop from fewer stores are more likely to be buyers who need to purchase everything from one of the stores in their network. Because of relative price dispersion, i.e., because equally expensive stores posts different prices, a buyer who is willing to purchase different goods at different stores in his network can pay less for the same basket of goods than a buyer who needs to purchase everything from just one of the stores in his network. Thus, households who shop from more stores tend to have a lower price index.

The theory makes other predictions about the dispersion in the prices paid by different households for the same basket of goods. Before reviewing these predictions, we look at the extent and sources of households' price index dispersion in the KNCP data. We normalize the price p_{ijmt} at which good j is traded in transaction i taking place in market m during quarter t by the average price of that good in that market and in that quarter. We then

decompose the price p_{ijmt} into a store component, a store-good component, and a transaction component. The store component is defined as the average of the (normalized) price of the goods sold by the store s_i where transaction i took place. The store-good component is defined as the difference between the average of the (normalized) price at which good j is sold at store s_i during quarter t and the store component. The transaction component is defined as the difference between the (normalized) price at which good j is sold in transaction i and the average of the (normalized) price at which good j is sold at store s_i during quarter t and the store component. We refer the reader to Kaplan and Menzio (2015) for details about this decomposition. For our purposes, it is important to notice that, roughly speaking, the store component and the store-good component of p_{ijmt} are “persistent” parts of the price, as they are measured as quarterly averages. The transaction component is a “transitory” part of the price, as it is measured as a deviation from a quarterly average.

We can decompose the households’ price indexes using the decomposition of the prices p_{ijmt} . Specifically, we express the price index of a household as the sum of a store component, a store-good component, and a transaction component. The store component of the price index is the (expenditure-weighted) average of the store component of all household transactions. The store-good component of the price index is the (expenditure-weighted) average of the store-good component of all household transactions. The transaction component of the price index is the (expenditure-weighted) average of the transaction component of all household transactions. Then, we decompose the variance of the households’ price indexes in market m and quarter t into the variance of the store component, the variance of the store-good component, the variance of the transaction component, and two covariance terms.

The average standard deviation of price indexes is 9%, and the average variance is 0.81%. The fraction of the variance due to differences in the store component is 42%, the fraction due to differences in the store-good component is 58%, and the fraction due to differences in the transitory store-good component is 19%. The covariance between the store component and the store-good component is -24% , while the covariance between the store component and the transaction component is 5%. Since ours is a theory of the persistent component of prices, we need to purge the price indexes from the transaction component. After doing so, we find that the average standard deviation of price indexes is 7.8%, and the average variance is 0.61%. The fraction of the variance due to differences in the store component is 55%, and the fraction due to differences in the persistent store-good component (plus the covariance term) is 45%.

Two features of price index dispersion are remarkable. First, the standard deviation of

price indexes (7.8%) is smaller than the standard deviation of prices (10%). Second, the fraction of the variance of price indexes due to differences in the store component (55%) is higher than the fraction of the variance of prices due to differences in the store component (35%). Conversely, the fraction of the variance of price indexes due to differences in the store-good component (together with the covariance term) is smaller than the variance of prices due to differences in the persistent part of the store-good component (45% vs 65%).

Our theory accounts surprisingly well for both features of price index dispersion, even though it was not designed or calibrated to do so. First, in the calibrated model, the standard deviation of price indexes is 6.6%, while the standard deviation of prices is 10%. Thus, the model can account, both qualitatively and quantitatively, for the fact that the dispersion of prices paid by households for the same basket of goods is lower than the dispersion of prices posted by sellers for the same good. Second, in the calibrated model, the fraction of the variance of price indexes due to differences in the store component is 71%, while the fraction of the variance of prices due to differences in the store component is 37%. Conversely, the fraction of the variance of price indexes due to the store-good component (together with the covariance) is 29%, while the fraction of the variance of prices due to the persistent part of the store-good component is 63%. Thus, the model can account qualitatively for the fact that differences in the store component are a larger source of dispersion of prices paid by households for the same basket of goods than they are a source of dispersion of prices posted by sellers for the same good.

The explanation for these phenomena provided by our theory is simple. Some sellers post a relatively high price for the first good and a relatively low price for the second good. Other equally expensive sellers do the opposite, as they post a relatively low price for the first good and a relatively high price for the second good. The variation across stores in the store-good component of prices contributes to a large fraction of the overall variance of posted prices. Now, let us turn to the buyers. Recall that buyers of type b have to purchase both goods from the same seller. In the price index of this type of buyer, the variation in the store-good component of prices washes out, as one of the prices they pay has a positive store-good component and the other has a negative store-good component. For this reason, the dispersion in price indexes across buyers is smaller than the dispersion in prices across sellers. Also for this reason, the variation in the store component of prices is more important for the dispersion of price indexes across buyers than for the dispersion of prices across sellers.

To summarize, our simple theory of pricing and shopping in the retail market provides an explanation for the extent and sources of dispersion in the prices posted by different sellers for the same good, for the buyers' return from shopping at an additional store, and

for the extent and sources of dispersion in the prices paid by different households for the same basket of goods. Thus, our theory responds to the call in Kaplan and Menzio (2015) for a model of the retail market that is consistent not only with the key facts about price dispersion but also with the key differences between the extent and sources of price and price index dispersion.

4 Conclusions

In this paper, we used the KNRS to measure the extent and sources of price dispersion, i.e. the dispersion in the price at which the same good is sold by different stores in the same market and in the same week. We found that a significant fraction of price dispersion is due to the fact that stores that are, on average, equally expensive choose to set persistently different prices for the same good. We labelled this phenomenon *relative price dispersion*. We then developed a theory of relative price dispersion in the context of the canonical model of price dispersion of Burdett and Judd (1983). According to our theory, relative price dispersion is an equilibrium manifestation of the sellers' attempt to discriminate between different buyers. In particular, an individual seller finds it optimal to charge asymmetric prices for different goods to discriminate between high-valuation buyers who need to purchase everything in the same location, and low-valuation buyers who are willing to purchase different goods at different locations. In equilibrium, for every seller that charges a relatively high price for one good, there must be another seller that is equally expensive on average but charges a relatively low price for the same good. We showed that our theory of relative price dispersion can not only account for the extent and sources of price dispersion, but also for the extent and sources of variation in the prices paid by different households for the same basket of goods. Thus, our theory offers a response to the call in Kaplan and Menzio (2015) for a model of the retail market that is consistent not only with the key facts about price dispersion but also with the key facts about the dispersion of price indexes across households.

Several extensions of our theory seem worthwhile. On the descriptive side, it would be interesting to combine our model of the retail market with a model of temporary price reductions (as Sobel 1984 or Aguirregabiria 1999). The resulting model would offer a truly comprehensive theory of price dispersion, in which price dispersion occurs because some stores are cheap and some are expensive, because equally expensive stores have different average prices for the same good, and because the same store has a different price for the same good on different days. As in this paper, the resulting theory could be tested using data on the dispersion in prices paid by different households. On the normative side, it

would be interesting to use the calibrated model to measure the extent of inefficiency in the retail market, identify which policies might be welfare improving, and which ones might exacerbate inefficiencies. Finally, it might be worthwhile extending the model presented in this paper to include a richer pattern of heterogeneity among buyers (i.e., more than two types of buyers), a richer set of goods (i.e., more than two goods), to endogenize the buyers' network of sellers, the buyers' decision to shop from multiple sellers or not, and the sellers' entry decision.

References

- AGUIAR, M., AND E. HURST (2005): "Consumption versus expenditure," *Journal of Political Economy*, 113(5), 919–948.
- (2007): "Lifecycle Prices and Production," *American Economic Review*, 97(5), 1533–1559.
- AGUIRREGABIRIA, V. (1999): "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, 66(2), 275–308.
- ALBRECHT, J., AND B. AXELL (1984): "An Equilibrium Model of Search Unemployment," *Journal of Political Economy*, 92(5), 824–840.
- ALBRECHT, J., F. POSTEL-VINAY, AND S. VROMAN (2013): "An Equilibrium Search Model of Synchronized Sales," *International Economic Review*, 54(2), 473–493.
- ALESSANDRIA, G., J. P. KABOSKI, AND V. MIDRIGAN (2010): "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, 100(5), 2304–2339.
- ATTANASIO, O. P., AND L. PISTAFERRI (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30(2), 3–28.
- BAUGHMAN, G., AND K. BURDETT (2015): "Multi-Product Search Equilibria," Mimeo, University of Pennsylvania.
- BELL, D. R., AND J. M. LATTIN (1998): "Shopping Behavior and Consumer Preference for Retail Price Format: Why Large Basket Shoppers Prefer EDLP," *Marketing Science*, 17(1), 66–88.
- BLUNDELL, R., AND I. PRESTON (1998): "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113(2), 603–640.
- BRYNJOLFSSON, E., AND M. D. SMITH (2000): "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 46(4), 563–585.
- BURDETT, K., AND K. L. JUDD (1983): "Equilibrium Price Dispersion," *Econometrica*, 51(4), 955–969.
- BUTTERS, G. R. (1977): "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 44(3), 465–491.
- CAMPBELL, J., AND B. EDEN (2014): "Rigid Prices: Evidence from U.S. Scanner Data," *International Economic Review*, 55(2), 423–442.

- CHEN, Z., AND P. REY (2012): “Loss Leading as an Exploitative Practice,” *American Economic Review*, 102, 3462–3482.
- CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): “Cyclic Pricing by a Durable Goods Monop-
olist,” *Quarterly Journal of Economics*, 99(3), 489–505.
- DAHLBY, B., AND D. S. WEST (1986): “Price Dispersion in an Automobile Insurance Market,”
Journal of Political Economy, 94(2), 418–438.
- ELLICKSON, P. B., AND S. MISRA (2008): “Supermarket Pricing Strategies,” *Marketing Science*,
27(5), 811–828.
- ELLISON, S. F., C. M. SNYDER, AND H. ZHANG (2015): “Costs of Managerial Attention and
Activity as a Source of Sticky Prices: Structural Estimates from an Online Market,” Mimeo,
Massachusetts Institute of Technology.
- GALENIANOS, M., R. L. PACULA, AND N. PERSICO (2012): “A Search-Theoretic Model of the
Retail Market for Illicit Drugs,” *Review of Economic Studies*, 79(3), 1239–1269.
- GORODNICHENKO, Y., V. SHEREMIROV, AND O. TALAVERA (2016): “Price Setting in Online
Markets: Does IT Click?,” Mimeo, University of California at Berkeley.
- GOTTSCHALK, P., AND R. MOFFITT (1994): “The Growth of Earnings Instability in the U.S.
Labor Market,” *Brookings Papers on Economic Activity*, 25(2), 217–272.
- HONG, H., AND M. SHUM (2006): “Using Price Distributions to Estimate Search Costs,” *RAND
Journal of Economics*, 37(2), 257–275.
- HOTELLING, H. (1929): “Stability in Competition,” *Economic Journal*, 39, 41–57.
- KAPLAN, G. (2012): “Inequality and the Life Cycle,” *Quantitative Economics*, 3, 471–525.
- KAPLAN, G., AND G. MENZIO (2015): “The Morphology of Price Dispersion,” *International
Economic Review*, 56(4), 1165–1206.
- (2016): “Shopping Externalities and Self-Fulfilling Unemployment Fluctuations,” *Journal
of Political Economy*, 124(3), 771–825.
- LACH, S. (2002): “Existence And Persistence Of Price Dispersion: An Empirical Analysis,” *Review
of Economics and Statistics*, 84(3), 433–444.
- LAL, R., AND C. MATUTES (1989): “Price Competition in Multimarket Duopolies,” *RAND Jour-
nal of Economics*, 20(4), 516–537.
- (1994): “Retail Pricing and Advertising Strategies,” *Journal of Business*, 67(3), 345–370.
- LEWIS, M. (2006): “Price Dispersion and Competition With Differentiated Sellers,” *Journal of
Industrial Economics*, 56, 654–678.
- MCAFEE, R. P. (1995): “Multiproduct Equilibrium Price Dispersion,” *Journal of Economic The-
ory*, 67, 83–105.
- MENZIO, G., AND N. TRACHTER (2015a): “Equilibrium Price Dispersion Across and Within
Stores,” PIER Working Paper 15-003.

- (2015b): “Equilibrium Price Dispersion with Sequential Search,” *Journal of Economic Theory*, 160(6), 188–215.
- NAKAMURA, E., AND J. STEINSSON (2008): “Pass-Through in Retail and Wholesale,” *American Economic Review*, 98(2), 430–437.
- PRATT, J. W., D. A. WISE, AND R. ZECKHAUSER (1979): “Price Differences in Almost Competitive Markets,” *Quarterly Journal of Economics*, 93(2), 189–211.
- REINGANUM, J. F. (1979): “A Simple Model of Equilibrium Price Dispersion,” *Journal of Political Economy*, 87(4), 851–858.
- RHODES, A. (2015): “Multiproduct Retailing,” *Review of Economic Studies*, 82(1), 333–359.
- RHODES, A., AND J. ZHOU (2015): “Consumer Search and Retail Market Structure,” Mimeo, New York University.
- ROB, R. (1985): “Equilibrium Price Distributions,” *Review of Economic Studies*, 52(3), 487–504.
- SOBEL, J. (1984): “The Timing of Sales,” *Review of Economic Studies*, 51(3), 353–368.
- SORENSEN, A. T. (2000): “Equilibrium Price Dispersion in Retail Markets for Prescription Drugs,” *Journal of Political Economy*, 108(4), 833–850.
- STAHL, D. O. (1989): “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, 79(4), 700–712.
- STIGLER, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69(3), 213–225.
- STIGLITZ, J. E. (1987): “Competition and the Number of Firms in a Market: Are Duopolies More Competitive than Atomistic Markets?,” *Journal of Political Economy*, 95(5), 1041–1061.
- VARIAN, H. R. (1980): “A Model of Sales,” *American Economic Review*, 70(4), 651–659.
- WOODWARD, S., AND R. E. HALL (2012): “Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage-Market Evidence,” *American Economic Review*, 102(7), 3249–3276.
- ZHOU, J. (2014): “Multiproduct Search and the Joint Search Effect,” *American Economic Review*, 104(9), 2918–2939.

Appendix

A Robustness

In this appendix, we report estimates and variance decompositions under: (i) alternative specifications of the inclusion criteria for the data; (ii) alternative specifications of the statistical model; (iii) alternative sets of products; (iv) alternative markets; (v) alternative definitions of a market. The main take-away is that relative price dispersion is a very robust feature of the data.

A.1 Sample Selection

Our baseline selection criteria required that a minimum of $N_1 = 250$ of the 1,000 goods in our sample be sold at a given store in a given week for that store/week to be included in the estimation sample. Table 5 reports variance decompositions for $N_1 \in \{50, 500\}$ and shows that the results are not sensitive to this particular threshold. Our baseline selection criteria also required that a minimum of $N_2 = 50$ stores have positive sales for a given good in a given week for that good/week to be included in the sample. Table 5 also reports variance decompositions for $N_2 \in \{25, 100\}$ and shows that the results are also not sensitive to this threshold. Thus, it is unlikely that relative price dispersion is a statistical artifact of small samples and/or insufficient overlap of goods across stores. This is important because for some of the sets of UPCs considered below, we are required to set $N_1 = 50$ and $N_2 = 25$ in order to have sufficient overlap for reliable estimation.

A.2 Statistical Model

In our baseline specification, we modeled the transitory part of the store-good component as an MA(1), which implicitly assigns all price changes with a duration greater than one week to the persistent component. Since the transitory component is intended to capture the effects of temporary sales, the reader may be concerned that, if some sales last more than one week, our baseline specification may be interpreting some sales-induced price variation as relative price dispersion. To show that this is not the case, Table 6 reports the variance decomposition when we model the temporary component as an MA(5) or MA(10), thus allowing the transitory component to capture sales that potentially last up to 10 weeks. The decomposition is barely affected — the persistent parts of the store-good component still account for at least one-third of the variance of the store-good component.

Table 5: Robustness to sample criteria

	$N_1 = 50$		$N_1 = 500$		$N_2 = 25$		$N_2 = 100$	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
<u>Store</u>								
Transitory	0.017	5.0	0.011	3.6	0.011	3.2	0.011	3.3
Fixed plus Pers.	0.075	95.0	0.055	96.4	0.058	96.8	0.063	96.7
Total Store	0.077	19.0	0.056	13.6	0.059	15.3	0.064	16.8
<u>Store-good</u>								
Transitory	0.126	63.3	0.113	65.4	0.111	63.8	0.114	65.0
Fixed plus Pers.	0.096	36.7	0.082	34.6	0.084	36.2	0.084	35.0
Total Store-good	0.158	81.0	0.140	86.4	0.139	84.7	0.141	83.2
<u>Total</u>	0.176	100.0	0.151	100.0	0.151	100.0	0.155	100.0

Notes: This table presents a robustness exercise comparing our baseline results to results obtained using alternative cutoffs for required numbers of observations.

The reader may also be concerned that modeling temporary sales as an MA process of any order may fail to capture salient features of sales dynamics and that this may lead us to understate the importance of sales. To show that this is not the case, we consider two alternative approaches to modeling sales. First, we recognize that the cross-sectional distribution of prices induced by periodic sales is likely to be negatively skewed, which may help in identifying the component of price dispersion that is due to sales. To account for this possibility we assume that the transitory innovation to the store-good component, $\varepsilon_{j_s,t}^z$, is drawn from a skewed distribution, whose skewness we estimate alongside the other parameters of the model.¹⁵ As expected, the estimated skewness is mildly negative (coefficient of skewness = -0.6), which is consistent with our prior view of sales. However, the decomposition of the variance of prices, shown in Table 6, is not affected.

Second, we depart from the assumption of an MA process for the transitory store-good component and replace it with a process that more explicitly resembles temporary sales. In this model, the transitory store-good component is modeled as a 2-point process. In a given week, there is a probability ϕ that each good is on sale, and sales are independently distributed across goods and over time. If a good is on sale, then it is discounted from its regular price by a fraction δ . We assume that all sales last exactly one week, and the day that a good goes on sale is uniformly distributed within the week. This means that each sale will affect the observed price of the good in two adjacent weeks, so the auto-covariance of prices

¹⁵Note that our estimation procedure does not require distributional assumptions on the innovations. Our baseline specification and procedure uses only second moments of the price data and only estimates the distribution of price innovations up to their second moments. In order to achieve identification of third moments, we include joint third moments of prices.

Table 6: Robustness to statistical model

	MA(5)		MA(10)		Skewed MA(1)		Sales	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
<u>Store-good</u>								
Transitory	0.114	65.6	0.114	66.1	0.113	64.1	0.098	57.7
Fixed plus Pers.	0.082	34.4	0.082	33.9	0.084	35.9	0.084	42.3
Total Store-good	0.141	100.0	0.141	100.0	0.141	100.0	0.141	100.0

Notes: This table shows variance decompositions for the store-good component after replacing the MA(1) process with: (i) an MA(5); (ii) an MA(10); (iii) an MA(1) that allows skewness in the disturbances; and (iv) an explicit model of sales, as described in the main text.

is impacted at both the zero and first lag (as with an MA(1) process). Since we work with normalized prices, there is an additional restriction that the mean value of the transitory component, after accounting for sales, is zero. The estimated weekly sales probability is $\phi = 4.65\%$, and the corresponding average discount is $\delta = 52\%$.¹⁶ The associated variance decomposition, show in Table 6, reveals that the persistent components of the store-good variance are even larger than in the baseline. We conclude that relative price dispersion is not driven by temporary sales and is a distinct feature of price distributions.

A.3 Products

Our baseline set of goods is the 1,000 most-commonly purchased products in Minnesota in the first quarter of 2010. We now show that relative price dispersion is not specific to this set of goods but rather is a robust phenomenon that is present among samples of products chosen in a broad variety of ways. The analysis also serves the purpose of rejecting some “banal” explanations for relative price dispersion, such as managerial inattention, store-good cost differentials, and different styles of shelf management.

A.3.1 Frequency of Purchase

Our baseline procedure weights each good equally when constructing the good-time means and the store components. In Table 7, we report the variance decomposition when we use quantity weights to construct the good-time means and store components. The decomposition is barely affected by this change.

¹⁶As with the skewed MA process, the sales process requires joint third moments of prices to be included in the GMM objective in order to achieve identification. We have also explored richer specifications in which the sales discount, δ , is itself a (possibly negatively skewed) random variable, and none of our main findings are affected.

Our baseline sample comprises only goods that are purchased very frequently. We examine whether relative price dispersion is a feature of the data for less frequently purchased goods. To do this, we select a sample of the 1,000 goods ranked 9,001 to 10,000 in terms of their frequency of purchase in Minnesota in the first quarter of 2010. This choice is motivated by our desire to select substantially less-commonly purchased goods than in our baseline sample, while still satisfying the requirement that the goods are sufficiently commonly purchased so that there is enough overlap across stores and enough continuity in weekly sales to meet our two inclusion criteria. The types of goods in this alternative sample, shown in Table 11 in Appendix B, are quite different from those in the baseline sample. However, the variance decomposition for this set of goods, shown in Table 7 (labelled “UPC-alt”), is extremely similar to the baseline.

Lastly, we selected a different sample of goods based on frequently purchased goods *nationwide* in the first quarter of 2010 rather than frequently purchased goods in Minnesota. Selecting a sample in this way is useful for when we extend our analysis to other parts of the country below. To construct this sample, we created two lists of the most commonly purchased UPCs, one based on quantity and one based on revenue. We then selected the 1,463 goods that appear in either list. The decomposition for this set of 1,463 goods is also shown in Table 7 (labelled “UPC-national”). For this set of goods, the store component accounts for slightly more of the overall price variation, but the persistent components of prices account for even more of the variance of the store-good component. Hence, relative price dispersion is larger in this set of goods than in the baseline.

A.3.2 High-Price and Low-Price Goods

A simple explanation for relative price dispersion is managerial inattention (see, e.g., Ellison, Snyder, and Zhang 2015). Equally expensive stores may set persistently different prices for the same good because managers choose to not pay much attention to the price of low-ticket items.¹⁷ This potential explanation for relative price dispersion motivates us to decompose price dispersion for low- and high-price goods separately. We divide our baseline sample of 1000 UPCs according to their average unit price. The low-price subsample of 430 UPCs has a median unit average price of 99 cents, a 5th percentile of 39 cents and a 95th percentile of \$1.79; the high-price subsample of 315 UPCs has a median unit average price of \$3.59 cents, a 5th percentile of \$2.39 and a 95th percentile of \$6.99. The variance decompositions for these two subsamples are shown in Table 7. The low-price subsample features more relative price dispersion than the full sample: The store-good component accounts for 79%

¹⁷We thank Stephan Seiler for suggesting this hypothesis.

Table 7: Robustness

	Baseline		Weighted		UPC-alt		UPC-national	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
<u>Store</u>								
Transitory	0.011	3.2	0.019	21.0	0.006	1.4	0.011	2.2
Fixed plus Pers.	0.059	96.8	0.037	79.0	0.054	98.6	0.072	97.8
Total Store	0.060	15.5	0.041	6.5	0.055	18.7	0.073	18.9
<u>Store-good</u>								
Transitory	0.113	64.1	0.124	62.6	0.082	51.1	0.119	61.0
Fixed plus Pers.	0.084	35.9	0.096	37.4	0.080	48.9	0.095	39.0
Total Store-good	0.141	84.5	0.156	93.5	0.114	81.3	0.152	81.1
	Low price		High price		Low durability		High durability	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
<u>Store</u>								
Transitory	0.024	8.7	0.025	15.6	0.013	4.0	0.027	27.9
Fixed plus Pers.	0.078	91.3	0.059	84.4	0.062	96.0	0.043	72.1
Total Store	0.082	20.6	0.065	15.9	0.063	19.3	0.051	19.4
<u>Store-good</u>								
Transitory	0.122	57.4	0.130	77.0	0.103	64.0	0.077	55.8
Fixed plus Pers.	0.105	42.6	0.071	23.0	0.077	36.0	0.069	44.2
Total Store-good	0.161	79.4	0.148	84.1	0.129	80.7	0.103	80.6
	Unilever		Coca-Cola		State: MN		County: Hennepin	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
<u>Store</u>								
Transitory	0.035	27.4	0.030	15.5	0.011	2.5	0.015	6.2
Fixed plus Pers.	0.058	72.6	0.070	84.5	0.070	97.5	0.058	93.8
Total Store	0.068	21.3	0.076	26.2	0.071	17.6	0.060	12.5
<u>Store-good</u>								
Transitory	0.101	60.9	0.106	68.9	0.120	60.9	0.128	64.4
Fixed plus Pers.	0.081	39.1	0.071	31.1	0.096	39.1	0.095	35.6
Total Store-good	0.130	78.7	0.127	73.8	0.154	82.4	0.159	87.5

Notes: This table presents a robustness exercise comparing our baseline results to results obtained using alternative specifications: quantity weighting in constructing the store and store-good components, alternative samples of UPCs (UPC-alternative and UPC-national), the low- and high-price samples, the low- and high-durability samples, the Unilever and Coca-Cola samples, and alternative definitions of a market (state of Minnesota and Hennepin County). Tables 8, and 11-18 in Appendix B illustrate the diversity in product groups across the alternative samples.

of the overall variance of prices, of which the persistent components account for 43%. The high-price subsample features less relative price dispersion than the full sample, but relative price dispersion is still a substantial fraction of overall price dispersion. Hence, relative price dispersion is not only a feature of low-price, low-revenue goods and thus is unlikely to be entirely due to managerial inattention.

A.3.3 Products from a Single Distributor

Another possible explanation for relative price dispersion is that equally expensive stores set persistently different prices for the same good because they have better or worse relationships (and, hence, are charged lower or higher prices) with the wholesaler.¹⁸ This potential explanation motivates us to decompose price dispersion for a subset of products produced and distributed by a single wholesaler. Indeed, if relative price dispersion is caused by different retailer-wholesaler relationships, relative price dispersion should be absorbed by the store component when we restrict attention to products from a single wholesaler.

We consider two subsamples of goods. In the first subsample, there are only products from Coca-Cola. In the second subsample, there are only products from Unilever. The 3,608 UPCs in our Coca-Cola subsample are primarily various types of beverages. The 10,866 UPCs in our Unilever subsample come from a variety of product groups; “Hair Care” is the product group with the largest fraction of UPCs (32%), followed by “Personal Soap and Bath Additives” (13%), “Deodorant” (12%) and “Skin Care Preparations” (10%).

The variance decompositions for these two subsamples of goods is shown in the bottom row of Table 7. For both samples, the overall degree of price dispersion is very similar to the degree of price dispersion in our baseline sample. However, the fraction of variation that is due to the store component is somewhat larger – 21% for Unilever and 26% for Coca-Cola, compared with 16% for the baseline. This is consistent with the hypothesis that some part of price dispersion is due to different relationships between particular stores and particular distributors. However, for both of these distributors, the vast majority of price dispersion is due to the store-good component, and, of this, the persistent parts account for 39% (Unilever) and 31% (Coca-Cola). Thus, relative price dispersion exists even when only considering goods from the same distributor and so is not only driven by heterogeneity in distributional relationships.

¹⁸We thank Matthew Gentzkow for suggesting this hypothesis.

A.3.4 High-Durability and Low-Durability Products

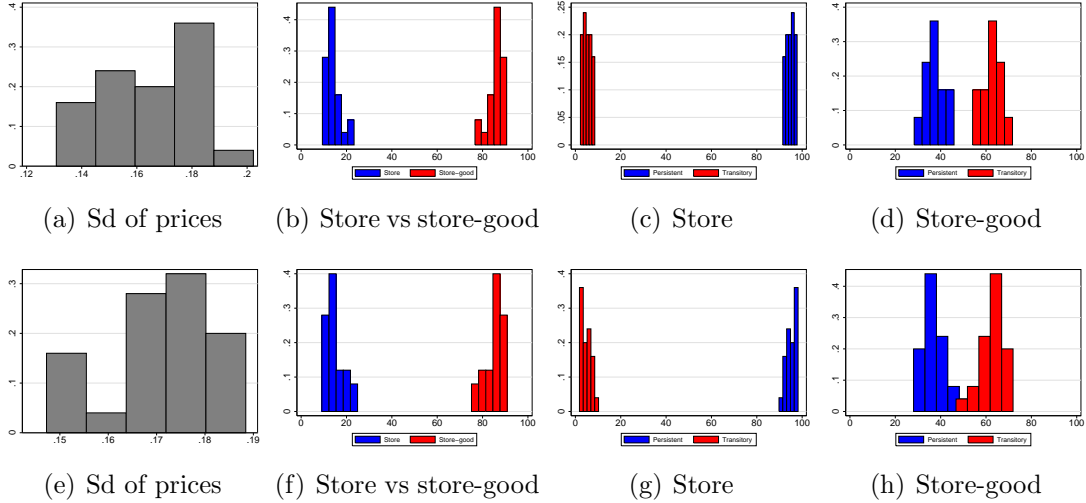
Another natural explanation for relative price dispersion is shelf management. Some stores may keep perishable goods on their shelves for longer and, for this reason, sell them at systematically lower prices, while other stores may remove perishable goods sooner and, for this reason, sell them at systematically higher prices. This observation motivates us to decompose price dispersion separately for two subsamples of goods: low-durability goods (i.e., perishable goods) and high-durability goods.¹⁹ The variance decompositions for these two subsamples are shown in Table 7. Even though the two subsamples contain very different sets of products, the overall decomposition of price dispersion is quite similar. For both subsamples, the store component accounts for approximately 20% and the store-good component for 80% of the cross-sectional variance of prices. For both subsamples, the transitory part accounts for roughly two-thirds and the persistent part for roughly one-third of the cross-sectional variance of the store-good component of prices. These findings suggest that relative price dispersion is unlikely to be a phenomenon caused by different styles of shelf management for perishable goods. Indeed, relative price dispersion turns out to be slightly more important in the subsample of goods that are less perishable.

A.4 Markets

So far our analysis has focused on a single geographic region — the Minneapolis-St Paul DMA. Here, we show that none of our results are specific to this level of geographic aggregation or this part of the country. First, we consider alternative levels of geographic aggregation for the definition of a market. In Table 7, we report the variance decomposition when we use a broader definition of market (the state of Minnesota) and a narrower definition of a market (Hennepin County, which is contained in the Minneapolis-St Paul DMA). All our findings are robust to switching to either of these alternative levels of aggregation.

Second, we extend our analysis to the whole of the United States in order to verify that our findings are not specific to Minneapolis-St Paul. We present results both at the level of a DMA and the county level. For each level of geographic aggregation, we selected the 25 largest areas by revenue and repeated the estimation for each market, using the same set of 1,463 UPCs. As described earlier, this set of UPCs was chosen to reflect UPCs that are

¹⁹We thank Boyan Jovanovic for suggesting we measure relative price dispersion for low- and high-durability products. We also thank George Alessandria for sharing the durability indexes constructed in Alessandria, Kaboski, and Midrigan (2010). We merge this index with the Nielsen database at the product module level by comparing descriptions of products. We define low-durability goods as those with a durability index of less than 2 months, and high-durability goods as those with a durability index of more than 140 months.



Notes: These histograms present a robustness exercise by looking at how our results for price dispersion and the variance decompositions vary across geographic regions across the country. The top row considers DMAs across the US, and the bottom counties across the US.

Figure 7: Price dispersion and variance decompositions across geographic areas

commonly purchased nationwide.

Figure 7(a) displays a histogram of the standard deviation of prices in each of the 25 DMAs. The corresponding variance decomposition between the store and store-good components is shown in Figure 7(b), and the fraction of the variance of each component that is due to transitory versus persistent factors is shown in Figures 7(c) and 7(d), respectively. The analogous statistics are displayed for the 25 counties in the bottom row of Figure 7. Figure 7 shows very clearly that our findings are not unique to any one particular region but instead are a general feature of price dynamics and distributions. For all geographic areas, virtually all of the variance of prices is due to the store-good component rather than the store component, and a substantial part of the variance of the store-good component (between one-third and one-half) is very persistent in nature.

B Additional Tables

Table 8: Share of UPCs across Product Groups: Baseline

Product Group	Percent		
YOGURT	10.7	DOUGH PRODUCTS	0.5
CARBONATED BEVERAGES	9.3	FRUIT - DRIED	0.5
FRESH PRODUCE	6	JUICES, DRINKS-FROZEN	0.5
BREAD AND BAKED GOODS	5.4	SALAD DRESSINGS, MAYO, TOPPINGS	0.5
PIZZA/SNACKS/HORS D'OEUVRES-FRZN	4.4	SUGAR, SWEETENERS	0.5
MILK	3.6	BOOKS AND MAGAZINES	0.4
VEGETABLES - CANNED	3.4	DESSERTS, GELATINS, SYRUP	0.4
SOFT DRINKS-NON-CARBONATED	3.3	ICE CREAM, NOVELTIES	0.4
SOUP	3.3	ORAL HYGIENE	0.4
CANDY	3.2	BAKING SUPPLIES	0.3
CEREAL	3	BREAKFAST FOODS-FROZEN	0.3
FRESH MEAT	3	FRESHENERS AND DEODORIZERS	0.3
SNACKS	3	ICE	0.3
CHEESE	2.9	JAMS, JELLIES, SPREADS	0.3
PAPER PRODUCTS	2.8	BABY FOOD	0.2
BREAKFAST FOOD	2.3	BAKING MIXES	0.2
CRACKERS	2.1	DESSERTS/FRUITS/TOPPINGS-FROZEN	0.2
DRESSINGS/SALADS/PREP FOODS-DELI	1.8	FRUIT - CANNED	0.2
PREPARED FOOD-DRY MIXES	1.8	HOUSEHOLD CLEANERS	0.2
PASTA	1.7	KITCHEN GADGETS	0.2
EGGS	1.6	PACKAGED MILK AND MODIFIERS	0.2
JUICE, DRINKS - CANNED, BOTTLED	1.6	WRAPPING MATERIALS AND BAGS	0.2
COOKIES	1.3	AUTOMOTIVE	0.1
COT CHEESE, SOUR CREAM, TOPPINGS	1.3	COFFEE	0.1
BUTTER AND MARGARINE	1.2	DISPOSABLE DIAPERS	0.1
PREPARED FOODS-FROZEN	1.2	FLOUR	0.1
CONDIMENTS, GRAVIES, AND SAUCES	1.1	HARDWARE, TOOLS	0.1
PREPARED FOOD-READY-TO-SERVE	1.1	HOUSEHOLD SUPPLIES	0.1
VEGETABLES-FROZEN	1.1	LIGHT BULBS, ELECTRIC GOODS	0.1
PACKAGED MEATS-DELI	1	PICKLES, OLIVES, AND RELISH	0.1
GUM	0.8	SHORTENING, OIL	0.1
SEAFOOD - CANNED	0.7	SNACKS, SPREADS, DIPS-DAIRY	0.1
TOBACCO AND ACCESSORIES	0.6	SPICES, SEASONING, EXTRACTS	0.1
		UNPREP MEAT/POULTRY/SEAFOOD-FRZN	0.1

Table 9: Share of UPCs across Departments: Baseline

	Baseline	UPC Alt	UPC National	Coca-Cola	Unilever	Low durab.	High durab.	Low price	High price
ALCOHOLIC BEVERAGES	0%	0%	6%	0%	0%	0%	13%	0%	0%
DAIRY	22%	7%	12%	13%	2%	35%	0%	25%	11%
DELI	2%	1%	2%	0%	0%	0%	0%	1%	3%
DRY GROCERY	52%	49%	53%	87%	12%	4%	0%	58%	49%
FRESH PRODUCE	6%	1%	5%	0%	0%	9%	0%	4%	6%
FROZEN FOODS	9%	12%	4%	0%	8%	32%	0%	7%	11%
GENERAL MERCHANDISE	1%	4%	2%	0%	1%	0%	82%	1%	1%
HEALTH AND BEAUTY	0%	12%	2%	0%	62%	0%	4%	0%	1%
MEAT	4%	3%	3%	0%	0%	20%	0%	1%	10%
NON-FOOD	4%	11%	11%	0%	14%	0%	0%	3%	8%
TOTAL	1000	1000	1463	3608	10917	12301	32989	430	315

Table 10: Parameter estimates

Models	Store component				
	ρ^y	$\theta_{y,1}$	$Var(\alpha^y)$	$Var(\eta^y)$	$Var(\varepsilon^y)$
Baseline	0.982613573	0	0.002793056	2.46E-05	0.000116634
State	0.983490975	0.038776	0.002663621	2.25E-05	0.000130635
County	0.991166659	0.217979	0.001667647	3.04E-05	0.000214859
$N_1 = 50$	0.987028926	0.086678	0.005004708	1.47E-05	0.00028866
$N_1 = 500$	0.990499883	0	0.002280235	1.32E-05	0.000110802
$N_2 = 25$	0.983067304	0	0.002687951	2.39E-05	0.000112249
$N_2 = 100$	0.983882217	0.006013	0.003159055	2.40E-05	0.000131709
Quant weighted	0.98573063	0.154354	0.00103578	8.43E-06	0.000346942
UPC alt	0.989435059	0.039211	0.002554465	8.57E-06	4.07E-05
UPC national	0.868614266	0	0.005151485	2.49E-05	0.000119115
Low price	0.980758409	0.131075	0.00555965	2.16E-05	0.000576657
High price	0.980959918	0.06653	0.003284731	8.60E-06	0.000645019
Low durability	0.962261334	0	0.003505668	2.45E-05	0.000160791
High durability	0.985078092	0.408077	0.000928235	2.73E-05	0.000614076
Unilever	0.984127103	0.07472	0.002913701	1.28E-05	0.001244086
Coca-cola	0.995571178	0.0427	0.000230254	4.10E-05	0.000892025

Models	Store-good component				
	ρ^z	$\theta_{z,1}$	$Var(\alpha^z)$	$Var(\eta^z)$	$Var(\varepsilon^z)$
Baseline	0.965159943	0.025611	0.003273669	0.000262241	0.012660499
State	0.965000278	0.042116	0.003592662	0.00027266	0.013102838
County	0.964487279	0.034764	0.00267737	0.000245359	0.012897066
$N_1 = 50$	0.969583498	0.03929	0.004748385	0.000264688	0.015791481
$N_1 = 500$	0.964557595	0.016314	0.003151334	0.000252595	0.012829199
$N_2 = 25$	0.965167539	0.026807	0.003236118	0.000258127	0.012361954
$N_2 = 100$	0.964220785	0.032336	0.003191119	0.000266847	0.012973907
Identity weights	0.96724252	0.030042	0.003211072	0.000246603	0.012723392
MA(5)	0.970311168	1	0.003151381	0.000213223	0.006327444
MA(10)	0.980000019	0.115018	0.003150801	0.000204187	0.006453931
Skewed MA(1)	0.965159943	0.025611	0.003273669	0.000262241	0.012660499
Quant weighted	0.966668046	0.046261	0.004811795	0.000285567	0.015285712
UPC alt	0.967386315	0.241865	0.003009445	0.000216338	0.006292217
UPC national	0.967495404	0.073928	0.004770164	0.000269885	0.013976282
Low price	0.968290143	0.058805	0.005416631	0.000350237	0.014827101
High price	0.97228473	0	0.002602982	0.000133687	0.01691263
Low durability	0.966802325	0.056946	0.002670222	0.000216992	0.010632362
High durability	0.954017962	0.071649	0.001431966	0.000294292	0.005906032
Unilever	0.973958285	0.151937	0.002962665	0.000187877	0.01006858
Coca-cola	0.952179616	0.004329	0.002045257	0.000280396	0.01120506

	ρ^z	$Var(\alpha^z)$	$Var(\eta^z)$	sale prob	discount
Sales	0.966006219	0.003260204	0.000251805	0.046535283	-0.520418332

Notes: The baseline model is estimated on a baseline sample of UPCs using data for the Minneapolis-St Paul designated market area. The following rows present results for alternative specifications discussed in the text: defining the market as the state or county, alternative cutoffs for constructing the samples, quantity weighting in constructing the store components, the alternative selections of goods “UPC alt” and “UPC nationwide,” the low- and high-price samples, the low- and high-durability samples, and the Unilever and Coca-cola samples. The top table considers the store component and the bottom the store-good component. For the latter we further investigate an alternative GMM weighting matrix (identity weighting), as well as alternative specifications for the transitory variation: MA(5), MA(10), MA(1) with skewness in disturbances, and an explicit model of sales.

Table 11: Share of UPCs across Product Groups: Alternative

Product Group	Percent		
CANDY	5.31	SOFT DRINKS-NON-CARBONATED	0.8
PREPARED FOODS-FROZEN	3.5	BAKED GOODS-FROZEN	0.7
SNACKS	3.2	COUGH AND COLD REMEDIES	0.7
JUICE, DRINKS - CANNED, BOTTLED	2.8	FRESHENERS AND DEODORIZERS	0.7
PAPER PRODUCTS	2.8	FRUIT - CANNED	0.7
PACKAGED MEATS-DELI	2.6	HARDWARE, TOOLS	0.7
BREAD AND BAKED GOODS	2.3	HOUSEHOLD CLEANERS	0.7
CONDIMENTS, GRAVIES, AND SAUCES	2.3	HOUSEHOLD SUPPLIES	0.7
SOUP	2.3	LIGHT BULBS, ELECTRIC GOODS	0.7
ICE CREAM, NOVELTIES	2.2	TEA	0.7
PET FOOD	2.2	FRESH MEAT	0.6
DEODORANT	2.1	HAIR CARE	0.6
CHEESE	2	TOBACCO AND ACCESSORIES	0.6
CARBONATED BEVERAGES	1.9	WRAPPING MATERIALS AND BAGS	0.6
PIZZA/SNACKS/HORS D'OEUVRES-FRZN	1.9	COT CHEESE, SOUR CREAM, TOPPINGS	0.5
BABY FOOD	1.8	DOUGH PRODUCTS	0.5
COOKIES	1.8	GLASSWARE, TABLEWARE	0.5
VEGETABLES-FROZEN	1.8	JAMS, JELLIES, SPREADS	0.5
MEDICATIONS/REMEDIES/HEALTH AI	1.7	VITAMINS	0.5
BREAKFAST FOOD	1.6	DESSERTS/FRUITS/TOPPINGS-FROZEN	0.4
CEREAL	1.6	JUICES, DRINKS-FROZEN	0.4
DESSERTS, GELATINS, SYRUP	1.6	SNACKS, SPREADS, DIPS-DAIRY	0.4
DETERGENTS	1.6	SPICES, SEASONING, EXTRACTS	0.4
PREPARED FOOD-DRY MIXES	1.6	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	0.4
SALAD DRESSINGS, MAYO, TOPPINGS	1.6	BREAKFAST FOODS-FROZEN	0.3
VEGETABLES - CANNED	1.6	BUTTER AND MARGARINE	0.3
CRACKERS	1.5	GROOMING AIDS	0.3
FRUIT - DRIED	1.5	PASTA	0.3
SANITARY PROTECTION	1.5	PICKLES, OLIVES, AND RELISH	0.3
DRESSINGS/SALADS/PREP FOODS-DELI	1.4	PUDDING, DESSERTS-DAIRY	0.3
YOGURT	1.4	SEAFOOD - CANNED	0.3
MILK	1.3	VEGETABLES AND GRAINS - DRIED	0.3
ORAL HYGIENE	1.3	BATTERIES AND FLASHLIGHTS	0.2
PREPARED FOOD-READY-TO-SERVE	1.3	BOOKS AND MAGAZINES	0.2
COFFEE	1.2	COOKWARE	0.2
FRESH PRODUCE	1.2	PACKAGED MILK AND MODIFIERS	0.2
LAUNDRY SUPPLIES	1.1	PET CARE	0.2
BAKING SUPPLIES	1	AUTOMOTIVE	0.1
DISPOSABLE DIAPERS	1	ELECTRONICS, RECORDS, TAPES	0.1
BAKING MIXES	0.9	FEMININE HYGIENE	0.1
FIRST AID	0.9	FLOUR	0.1
SHAVING NEEDS	0.9	GRT CARDS/PARTY NEEDS/NOVELTIE	0.1
SKIN CARE PREPARATIONS	0.9	KITCHEN GADGETS	0.1
STATIONERY, SCHOOL SUPPLIES	0.9	SEASONAL	0.1
GUM	0.8	SHOE CARE	0.1
NUTS	0.8	SHORTENING, OIL	0.1
PERSONAL SOAP AND BATH ADDITIV	0.8	SUGAR, SWEETENERS	0.1
		TABLE SYRUPS, MOLASSES	0.1
		WINE	0.1

Table 12: Share of UPCs across Product Groups: National

Product Group	Percent		
CARBONATED BEVERAGES	7.59	DISPOSABLE DIAPERS	0.55
FRESH PRODUCE	5.06	TEA	0.55
JUICE, DRINKS - CANNED, BOTTLED	4.65	VEGETABLES-FROZEN	0.55
CANDY	4.58	WINE	0.55
PAPER PRODUCTS	4.38	ICE CREAM, NOVELTIES	0.48
BEER	4.24	PASTA	0.48
SNACKS	3.76	SALAD DRESSINGS, MAYO, TOPPINGS	0.48
SOFT DRINKS-NON-CARBONATED	3.69	SUGAR, SWEETENERS	0.48
BREAD AND BAKED GOODS	3.63	COUGH AND COLD REMEDIES	0.41
YOGURT	3.63	FRUIT - CANNED	0.41
CEREAL	3.49	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	0.41
TOBACCO AND ACCESSORIES	3.49	BATTERIES AND FLASHLIGHTS	0.34
PACKAGED MEATS-DELI	2.6	FRESH MEAT	0.34
PET FOOD	2.6	JAMS, JELLIES, SPREADS	0.34
MILK	2.53	SHAVING NEEDS	0.34
DRESSINGS/SALADS/PREP FOODS-DELI	2.26	WRAPPING MATERIALS AND BAGS	0.34
SOUP	2.26	BAKING SUPPLIES	0.27
CHEESE	2.05	DESSERTS, GELATINS, SYRUP	0.27
VEGETABLES - CANNED	1.92	HOUSEWARES, APPLIANCES	0.27
CONDIMENTS, GRAVIES, AND SAUCES	1.71	ICE	0.27
DETERGENTS	1.57	LAUNDRY SUPPLIES	0.27
PREPARED FOODS-FROZEN	1.37	SHORTENING, OIL	0.27
BABY FOOD	1.3	BAKING MIXES	0.21
BUTTER AND MARGARINE	1.23	BREAKFAST FOODS-FROZEN	0.21
CRACKERS	1.09	DOUGH PRODUCTS	0.21
GUM	1.09	ELECTRONICS, RECORDS, TAPES	0.21
LIQUOR	1.09	HOUSEHOLD SUPPLIES	0.21
PREPARED FOOD-READY-TO-SERVE	1.09	PICKLES, OLIVES, AND RELISH	0.21
EGGS	1.03	CHARCOAL, LOGS, ACCESSORIES	0.14
MEDICATIONS/REMEDIES/HEALTH AI	0.96	DESSERTS/FRUITS/TOPPINGS-FROZEN	0.14
PACKAGED MILK AND MODIFIERS	0.96	NUTS	0.14
PIZZA/SNACKS/HORS D'OEUVRES-FRZN	0.96	ORAL HYGIENE	0.14
PREPARED FOOD-DRY MIXES	0.89	SPICES, SEASONING, EXTRACTS	0.14
COT CHEESE, SOUR CREAM, TOPPINGS	0.82	DIET AIDS	0.07
SEAFOOD - CANNED	0.82	FLOUR	0.07
BREAKFAST FOOD	0.62	FRESHENERS AND DEODORIZERS	0.07
COFFEE	0.62	FRUIT - DRIED	0.07
COOKIES	0.62	HOUSEHOLD CLEANERS	0.07
BOOKS AND MAGAZINES	0.55	PHOTOGRAPHIC SUPPLIES	0.07
		SNACKS, SPREADS, DIPS-DAIRY	0.07
		STATIONERY, SCHOOL SUPPLIES	0.07
		VITAMINS	0.07

Table 13: Share of UPCs across Product Groups: Low Price

Product Group	Percent
YOGURT	21.63
CARBONATED BEVERAGES	10.23
VEGETABLES - CANNED	7.67
SOUP	6.51
SOFT DRINKS-NON-CARBONATED	6.05
CANDY	5.58
BREAD AND BAKED GOODS	4.42
FRESH PRODUCE	4.19
PASTA	3.72
PREPARED FOOD-DRY MIXES	3.02
PIZZA/SNACKS/HORS D'OEUVRES-FRZN	2.33
VEGETABLES-FROZEN	2.09
GUM	1.86
PAPER PRODUCTS	1.86
PREPARED FOOD-READY-TO-SERVE	1.86
EGGS	1.63
SEAFOOD - CANNED	1.63
JUICE, DRINKS - CANNED, BOTTLED	1.16
CHEESE	0.93
CONDIMENTS, GRAVIES, AND SAUCES	0.93
JUICES, DRINKS-FROZEN	0.93
PREPARED FOODS-FROZEN	0.93
DRESSINGS/SALADS/PREP FOODS-DELI	0.7
FRESHENERS AND DEODORIZERS	0.7
ICE	0.7
PACKAGED MEATS-DELI	0.7
BAKING MIXES	0.47
COT CHEESE, SOUR CREAM, TOPPINGS	0.47
DESSERTS/FRUITS/TOPPINGS-FROZEN	0.47
FRUIT - CANNED	0.47
KITCHEN GADGETS	0.47
ORAL HYGIENE	0.47
SUGAR, SWEETENERS	0.47
BAKING SUPPLIES	0.23
BUTTER AND MARGARINE	0.23
COOKIES	0.23
CRACKERS	0.23
DESSERTS, GELATINS, SYRUP	0.23
HARDWARE, TOOLS	0.23
LIGHT BULBS, ELECTRIC GOODS	0.23
MILK	0.23
PACKAGED MILK AND MODIFIERS	0.23
PICKLES, OLIVES, AND RELISH	0.23
SNACKS	0.23
SPICES, SEASONING, EXTRACTS	0.23

Table 14: Share of UPCs across Product Groups: High Price

Product Group	Percent
CARBONATED BEVERAGES	13.65
PIZZA/SNACKS/HORS D'OEUVRES-FRZN	8.89
FRESH MEAT	8.57
CEREAL	6.35
SNACKS	6.35
FRESH PRODUCE	6.03
PAPER PRODUCTS	5.08
BREAD AND BAKED GOODS	4.44
MILK	4.44
BREAKFAST FOOD	3.17
CRACKERS	3.17
COOKIES	2.54
DRESSINGS/SALADS/PREP FOODS-DELI	2.54
JUICE, DRINKS - CANNED, BOTTLED	2.54
SOFT DRINKS-NON-CARBONATED	2.22
TOBACCO AND ACCESSORIES	1.9
BUTTER AND MARGARINE	1.59
CANDY	1.59
COT CHEESE, SOUR CREAM, TOPPINGS	1.59
PACKAGED MEATS-DELI	1.59
CHEESE	1.27
EGGS	1.27
BOOKS AND MAGAZINES	0.95
ICE CREAM, NOVELTIES	0.95
SALAD DRESSINGS, MAYO, TOPPINGS	0.95
BABY FOOD	0.63
CONDIMENTS, GRAVIES, AND SAUCES	0.63
ORAL HYGIENE	0.63
PREPARED FOODS-FROZEN	0.63
YOGURT	0.63
BREAKFAST FOODS-FROZEN	0.32
COFFEE	0.32
DISPOSABLE DIAPERS	0.32
HOUSEHOLD CLEANERS	0.32
HOUSEHOLD SUPPLIES	0.32
SNACKS, SPREADS, DIPS-DAIRY	0.32
SUGAR, SWEETENERS	0.32
UNPREP MEAT/POULTRY/SEAFOOD-FRZN	0.32
VEGETABLES-FROZEN	0.32
WRAPPING MATERIALS AND BAGS	0.32

Table 15: Share of UPCs across Product Groups: Coca-cola

Product Group	Percent
CARBONATED BEVERAGES	56.26
JUICE, DRINKS - CANNED, BOTTLED	18.4
YOGURT	12.89
SOFT DRINKS-NON-CARBONATED	12.2
COFFEE	0.06
JUICES, DRINKS-FROZEN	0.06
VITAMINS	0.06
BUTTER AND MARGARINE	0.03
CHEESE	0.03
GLASSWARE, TABLEWARE	0.03

Table 16: Share of UPCs across Product Groups: Unilever

Product Group	Percent
HAIR CARE	31.83
PERSONAL SOAP AND BATH ADDITIV	13.55
DEODORANT	11.87
SKIN CARE PREPARATIONS	9.51
ICE CREAM, NOVELTIES	7.82
TEA	5.25
MEN'S TOILETRIES	2.78
SOUP	2.39
FRAGRANCES - WOMEN	2.1
GROOMING AIDS	1.95
YOGURT	1.39
SALAD DRESSINGS, MAYO, TOPPINGS	1.36
PREPARED FOOD-DRY MIXES	1.22
CONDIMENTS, GRAVIES, AND SAUCES	1.03
BUTTER AND MARGARINE	0.98
FIRST AID	0.95
SHAVING NEEDS	0.67
PAPER PRODUCTS	0.62
HOUSEWARES, APPLIANCES	0.58
DESSERTS/FRUITS/TOPPINGS-FROZEN	0.45
MEDICATIONS/REMEDIES/HEALTH AI	0.35
SPICES, SEASONING, EXTRACTS	0.32
COSMETICS	0.23
SANITARY PROTECTION	0.08
SHORTENING, OIL	0.08
SNACKS	0.08
SOFT DRINKS-NON-CARBONATED	0.08
JUICE, DRINKS - CANNED, BOTTLED	0.07
DRESSINGS/SALADS/PREP FOODS-DELI	0.06
DESSERTS, GELATINS, SYRUP	0.06
ETHNIC HABA	0.04
PREPARED FOOD-READY-TO-SERVE	0.04
BAKING SUPPLIES	0.03
VEGETABLES - CANNED	0.03
BABY NEEDS	0.02
BAKING MIXES	0.02
CANDY	0.02
AUTOMOTIVE	0.01
FRESHENERS AND DEODORIZERS	0.01
HOUSEHOLD CLEANERS	0.01
HOUSEHOLD SUPPLIES	0.01
JUICES, DRINKS-FROZEN	0.01
PACKAGED MILK AND MODIFIERS	0.01
SEWING NOTIONS	0.01
SNACKS, SPREADS, DIPS-DAIRY	0.01
VEGETABLES AND GRAINS - DRIED	0.01

Table 17: Share of UPCs across Product Groups: Low Durability

Product Group	Percent
ICE CREAM, NOVELTIES	31.52
PACKAGED MEATS-DELI	20.29
YOGURT	18.56
FRESH PRODUCE	9.28
MILK	8.59
CHEESE	4.11
BABY FOOD	3.16
COT CHEESE, SOUR CREAM, TOPPINGS	2.15
EGGS	1.78
BAKING SUPPLIES	0.28
DESSERTS/FRUITS/TOPPINGS-FROZEN	0.09
SNACKS, SPREADS, DIPS-DAIRY	0.02
SOFT DRINKS-NON-CARBONATED	0.02
CONDIMENTS, GRAVIES, AND SAUCES	0.02
DISPOSABLE DIAPERS	0.02
PAPER PRODUCTS	0.02
PET FOOD	0.02
PREPARED FOOD-READY-TO-SERVE	0.02
VITAMINS	0.02
CARBONATED BEVERAGES	0.01
CEREAL	0.01
GROOMING AIDS	0.01
HOUSEHOLD SUPPLIES	0.01
JAMS, JELLIES, SPREADS	0.01

Table 18: Share of UPCs across Product Groups: High Durability

Product Group	Percent
KITCHEN GADGETS	35.83
LIGHT BULBS, ELECTRIC GOODS	18.04
HOUSEWARES, APPLIANCES	13.32
LIQUOR	13.01
STATIONERY, SCHOOL SUPPLIES	10.08
PHOTOGRAPHIC SUPPLIES	4.73
BABY NEEDS	4.4
CHARCOAL, LOGS, ACCESSORIES	0.4
WINE	0.12
BEER	0.03
BATTERIES AND FLASHLIGHTS	0.01
JUICE, DRINKS - CANNED, BOTTLED	0.01
SOFT DRINKS-NON-CARBONATED	0.01
COOKWARE	0
DOUGH PRODUCTS	0
GRT CARDS/PARTY NEEDS/NOVELTIES	0
TOBACCO AND ACCESSORIES	0

C Proofs for Section 2

C.1 Proof of Lemma 1

(i) First, consider a seller posting prices (p_1, p_2) with $p_1 > u_b$ and $p_2 > u_b$. The seller's profit is zero, as there are no buyers willing to purchase a good at a price strictly greater than u_b . If the seller instead posts the prices (p'_1, p'_2) with $p'_1 = p'_2 = u_b$, it sells both goods to all captive buyers of type b and it attains a profit of at least $2\mu_b\alpha u_b > 0$. Therefore, a seller never finds it optimal to post prices (p_1, p_2) such that $p_1 > u_b$ and $p_2 > u_b$.

Consider a seller posting prices (p_1, p_2) with $p_1 > u_b$ and $p_2 \in (u_c, u_b]$. The seller attains a profit of

$$S(p_1, p_2) = \mu_b \left\{ \alpha + 2(1 - \alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} p_2. \quad (32)$$

The expression above is easy to understand. The seller is in the network of $\mu_b\alpha$ captive buyers of type b . A captive buyer of type b purchases good 2 from the seller with probability 1. The seller is also in the network of $2\mu_b(1 - \alpha)$ non-captive buyers of type b . A non-captive buyer of type b purchases good 2 from the seller with probability $1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2$, where $1 - \hat{G}(u_b + p_2)$ denotes the fraction of sellers that charge prices (p'_1, p'_2) such that $\min\{p'_1, u_b\} + \min\{p'_2, u_b\} > u_b + p_2$, and $\hat{\nu}(u_b + p_2)$ is the measure of sellers that charge prices (p'_1, p'_2) such that $\min\{p'_1, u_b\} + \min\{p'_2, u_b\} = u_b + p_2$. The seller is in the network of $\mu_c\alpha$ captive buyers of type c and of $2\mu_c(1 - \alpha)$ non-captive buyers of type c , but it does not trade with any of them as both of its prices are greater than u_c . If the seller instead posts the prices (p'_1, p'_2) with $p'_1 = u_b$ and $p'_2 = p_2$, it attains a profit of

$$\begin{aligned} S(p'_1, p'_2) &= \mu_b \left\{ \alpha + 2(1 - \alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} (u_b + p_2) \\ &> S(p_1, p_2). \end{aligned} \quad (33)$$

The strict inequality in (33) follows from the fact that, by lowering the price of good 1 to u_b , the seller trades both good 1 and good 2 to its customers of type b . Hence, a seller never finds it optimal to post prices (p_1, p_2) with $p_1 > u_b$ and $p_2 \in (u_c, u_b]$. For the same reason, a seller never finds it optimal to post prices (p_1, p_2) with $p_1 \in (u_c, u_b]$ and $p_2 > u_b$.

Finally, consider a seller posting prices (p_1, p_2) with $p_1 > u_b$ and $p_2 \in [0, u_c]$. The seller attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \left\{ \alpha + 2(1 - \alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} p_2 \\ &\quad + \mu_c \left\{ \alpha + 2(1 - \alpha) \left[1 - F_2(p_2) + \lambda_2(p_2)/2 \right] \right\} p_2. \end{aligned} \quad (34)$$

The expression in (34) is the same as in (32) with the addition of the term in the second line. This term represents the profit that the seller makes from trading good 2 to buyers of

type c . If the seller instead posts the prices (p'_1, p'_2) with $p'_1 = u_b$ and $p'_2 = p_2$, it attains a profit of

$$\begin{aligned} S(p'_1, p'_2) &= \mu_b \left\{ \alpha + 2(1 - \alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} (u_b + p_2) \\ &\quad + \mu_c \{ \alpha + 2(1 - \alpha) [1 - F_2(p_2) - \lambda_2(p_2)/2] \} p_2 \\ &> S(p_1, p_2) \end{aligned} \tag{35}$$

The strict inequality in (35) follows from the fact that, by lowering the price of good 1 to u_b , the seller trades both goods 1 and good 2 to its type b customers. Hence, a seller never finds it optimal to post prices (p_1, p_2) such that $p_1 > u_b$ and $p_2 \in [0, u_c]$. For the very same reason, a seller never finds it optimal to post prices (p_1, p_2) such that $p_1 \in [0, u_c]$ and $p_2 > u_b$.

(ii) Given part (i), it follows immediately that, if in equilibrium, a seller posts prices (p_1, p_2) with $p_1 + p_2 > u_b + u_c$, it must be the case that $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. ■

C.2 Proof of Lemma 3

(i) On the way to a contradiction suppose there is an equilibrium where G has a mass point at q_0 , i.e. $\nu(q_0) > 0$. First, notice that no seller finds it optimal to post $p_1 = p_2 = 0$, and hence the mass point cannot be at $q_0 = 0$. Second, notice that if, in equilibrium, a seller posts (p_1, p_2) with $p_1 + p_2 = q_0$, it must be posting $p_1 \in [0, u_b]$ and $p_2 \in [0, u_b]$. Therefore, this seller attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \{ \alpha + 2(1 - \alpha) [1 - G(q_0) + \nu(q_0)/2] \} q_0 \\ &\quad + \sum_{i=1}^2 \mu_c \{ \alpha + 2(1 - \alpha) [1 - F_i(p_i) + \lambda_i(p_i)/2] \} \mathbf{1}[p_i \leq u_c] p_i, \end{aligned} \tag{36}$$

where $\mathbf{1}[p_i \leq u_c]$ is the indicator function that takes the value 1 if $p_i \leq u_c$ and 0 otherwise. Suppose that the seller deviates and posts prices (p'_1, p'_2) with $0 \leq p'_1 = p_1 - \epsilon_1$, $0 \leq p'_2 = p_2 - \epsilon_2$, $\epsilon = \epsilon_1 + \epsilon_2$, where $\epsilon_1 \geq 0$, $\epsilon_2 \geq 0$ and $\epsilon > 0$ all arbitrarily small. Then, the seller attains a profit of

$$\begin{aligned} S(p'_1, p'_2) &= \mu_b \{ \alpha + 2(1 - \alpha) [1 - G(q_0 - \epsilon)] \} (q_0 - \epsilon) \\ &\quad + \sum_{i=1}^2 \mu_c \{ \alpha + 2(1 - \alpha) [1 - F_i(p_i - \epsilon_i) + \lambda_i(p_i - \epsilon_i)/2] \} \mathbf{1}[p_i \leq u_c] (p_i - \epsilon_i) \\ &> S(p_1, p_2), \end{aligned} \tag{37}$$

where the inequality follows from the fact that $G(q_0) - \nu(q_0)/2 - G(q_0 - \epsilon) \geq \nu(q_0)/2$ while ϵ , ϵ_1 and ϵ_2 are all arbitrarily small. Since $S(p'_1, p'_2) > S(p_1, p_2)$, there cannot be a mass point at q_0 .

(ii) On the way to a contradiction, suppose there is an equilibrium where F_1 has a mass point at $p_{1,0} \in (0, u_c]$. If in equilibrium a seller posts the price $p_1 = p_{1,0}$ for the first good, it

must post a price $p_2 \in [0, u_b]$ for the second good. This seller attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \{ \alpha + 2(1 - \alpha)[1 - G(p_1 + p_2)] \} (p_1 + p_2) \\ &\quad + \sum_{i=1}^2 \mu_c \{ \alpha + 2(1 - \alpha)[1 - F_i(p_i) + \lambda_i(p_i)/2] \} \mathbf{1}[p_i \leq u_c] p_i. \end{aligned} \quad (38)$$

If the seller deviates and posts prices (p'_1, p'_2) with $p'_1 = p_{1,0} - \epsilon$, $p'_2 = p_2$, for $\epsilon > 0$ arbitrarily small, it attains a profit of

$$\begin{aligned} S(p'_1, p'_2) &= \mu_b \{ \alpha + 2(1 - \alpha)[1 - G(p_{1,0} + p_2 - \epsilon)] \} (p_{1,0} + p_2 - \epsilon) \\ &\quad + \mu_c \{ \alpha + 2(1 - \alpha)[1 - F_1(p_{1,0} - \epsilon)] \} (p_{1,0} - \epsilon) \\ &\quad + \mu_c \{ \alpha + 2(1 - \alpha)[1 - F_2(p_2) + \lambda_2(p_2)/2] \} \mathbf{1}[p_2 \leq u_c] p_i \\ &\geq S(p_1, p_2), \end{aligned} \quad (39)$$

where the first inequality follows from the fact that the $F_1(p_{1,0}) - \lambda_1(p_{1,0})/2 - F_1(p_{1,0} - \epsilon) \geq \lambda_1(p_{1,0})/2$ and ϵ is arbitrarily small. Since $S(p'_1, p'_2) > S(p_1, p_2)$, there cannot be a mass point at $p_{1,0}$. ■

C.3 Proof of Proposition 1

We established part (i) in the main text. We have also established that, given the marginal distribution G in (14), a seller attains a profit of S^* for all $(p_1, p_2) \in R_1$ such that $p_1 + p_2 \in [q_\ell, q_h]$. In order to complete the proof of part (ii), all we need to do is find a condition under which a seller cannot attain a profit strictly greater than S^* by posting some off-equilibrium prices.

In Lemma 1, we proved that a seller never finds it optimal to post prices (p_1, p_2) with either p_1 , p_2 or both p_1 and p_2 strictly greater than u_b . If the seller posts prices $(p_1, p_2) \in R_1$ with $p_1 + p_2 \in (u_b + u_c, q_\ell)$, it attains a profit of

$$\begin{aligned} S_1(p_1 + p_2) &= \mu_b \{ \alpha + 2(1 - \alpha) \} (p_1 + p_2) \\ &< \mu_b \{ \alpha + 2(1 - \alpha) \} q_\ell = S^*, \end{aligned} \quad (40)$$

where the first line makes use of the fact that $G(p_1 + p_2) = 0$ and the second line makes use of the fact that $p_1 + p_2 < q_\ell$. Hence, a seller does not find it optimal to deviate from the equilibrium and post prices $(p_1, p_2) \in R_1$ with $p_1 + p_2 \in (u_b + u_c, q_\ell)$.

In Lemma 2, we proved that a seller never finds it optimal to post prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$ and both p_1 and p_2 greater than u_c and smaller than u_b . If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in [0, u_c]$ and $p_2 \in (u_c, u_b]$, it attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \{ \alpha + 2(1 - \alpha) \} (p_1 + p_2) + \mu_c \{ \alpha + 2(1 - \alpha) \} p_1 \\ &\leq \mu_b \{ \alpha + 2(1 - \alpha) \} (u_c + u_b) + \mu_c \{ \alpha + 2(1 - \alpha) \} u_c \\ &= S(u_c, u_b), \end{aligned} \quad (41)$$

where the first line makes use of $G(p_1 + p_2) = 0$ and $F(p_1) = 0$, and the second line makes use of $G(u_c + u_b) = 0$, $F(u_c) = 0$, $p_1 + p_2 \leq u_b + u_c$ and $p_1 \leq u_c$. The equilibrium profit S^* is greater than $S(u_c, u_b)$ if and only if

$$\frac{\mu_c}{\mu_b} \leq \frac{3\alpha - 2}{(2 - \alpha)u_c/u_b} - 1. \quad (42)$$

Hence, if and only if (42) holds a seller does not find it optimal to deviate from the equilibrium and post any prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in [0, u_c]$ and $p_2 \in (u_c, u_b]$. Similarly, condition (42) guarantees that a seller does not find it optimal to deviate from the equilibrium and post any prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in [0, u_c]$.

If the seller posts prices (p_1, p_2) with $p_1 \in [0, u_c]$ and $p_2 \in [0, u_c]$, it attains a profit of

$$\begin{aligned} S(p_1, p_2) &= (\mu_b + \mu_c) \{\alpha + 2(1 - \alpha)\} (p_1 + p_2) \\ &\leq (\mu_b + \mu_c) \{\alpha + 2(1 - \alpha)\} 2u_c \\ &= S(u_c, u_c), \end{aligned} \quad (43)$$

where the first line makes use of $G(p_1 + p_2) = 0$ and $F_i(p_i) = 0$ for $i = \{1, 2\}$, and the second line makes use of $G(2u_c) = 0$, $F_i(u_c) = 0$ for $i = \{1, 2\}$, and $p_1 + p_2 \leq 2u_c$. The equilibrium profit S^* is greater than $S(u_c, u_c)$ if and only if

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha}{(2 - \alpha)u_c/u_b} - 1. \quad (44)$$

Hence, if and only if (44) holds a seller does not find it optimal to deviate from the equilibrium and post any prices (p_1, p_2) with $p_1 \in [0, u_c]$ and $p_2 \in [0, u_c]$. Finally, notice that if condition (42) holds, so does condition (44). Therefore, a seller does not want to deviate from the equilibrium if and only if condition (42), which is the same as condition (16), is satisfied. ■

C.4 Proof of Lemma 5

It is easy to verify that in a Discrimination Equilibrium, $F_i(p)$ does not have a mass point at $p = 0$ for $i = \{1, 2\}$. Hence, $F_i(0) = 0$ and $F_i(p)$ is continuous over the interval $[0, u_c]$ for $i = \{1, 2\}$. If $F_1 \neq F_2$ for some $p \in [0, u_c]$, there exist prices p' and p'' , with $0 \leq p' < p'' \leq u_c$, for which $F_1(p') = F_2(p')$ and either $F_1(p) > F_2(p)$ or $F_1(p) < F_2(p)$ for all $p \in (p', p'')$. Without loss in generality, assume $F_1(p) > F_2(p)$ for all $p \in (p', p'')$. Any price $p \in (p', p'')$ for good 1 is posted by a seller in region R_2 , i.e. a seller with a basket price of $q \in (2u_c, u_b + u_c]$ and a price for good 2 of $p_2 \in (u_c, u_b]$. This seller attains a profit of $S_{21}(q, p)$. However, if

the seller inverts the prices of goods 1 and 2, it attains a profit of $S_{22}(q, p)$ which is strictly greater than $S_{21}(q, p)$ because $F_1(p) > F_2(p)$. In turn, $S_{21}(q, p) < S_{22}(q, p)$ implies that in equilibrium there are no sellers posting a price $p \in (p', p'')$ for good 1, i.e. $F_1(p) = F_1(p')$ for all $p \in (p', p'')$. Since $F_1(p) = F_1(p')$ for all $p \in (p', p'')$, $F_2(p) \geq F_2(p')$ for all $p \in (p', p'')$ and $F_1(p') = F_2(p')$, it follows that $F_2(p) \geq F_1(p)$ for all $p \in (p', p'')$. We have thus reached a contradiction. ■

C.5 Proof of Lemma 6

We need to prove that the seller's profit $S_2(q, p)$ is constant for all (q, p) such that $q \in [q_\ell, u_b + u_c]$ and $p \in [p_\ell, u_c]$. To this aim, it is useful to write the seller's profit $S_2(q, p)$ as

$$S_2(q, p) = B(q) + C(p), \quad (45)$$

where $B(q)$ and $C(p)$ are defined as

$$\begin{aligned} B(q) &= \mu_b [\alpha + 2(1 - \alpha)(1 - G(q))] q, \\ C(p) &= \mu_c [\alpha + 2(1 - \alpha)(1 - F(p))] p. \end{aligned}$$

In words, $B(q)$ is the profit that the seller attains from trades with the buyers of type b , and $A(p)$ is the profit that the seller attains from trades with the buyers of type c .

First, we show that the distribution $G(q)$ has full support over the interval $[q_\ell, u_b + u_c]$. On the way to a contradiction, suppose that $G(q)$ is such that $G'(q_0) > 0$ and $G(q_0) = G(q_1)$ for some q_0 and q_1 with $q_\ell < q_0 < q_1 \leq u_b + u_c$. Since $G(q) = G(q_0)$ for all $q \in [q_0, q_1]$, the function $B(q)$ is strictly increasing in q for all $q \in [q_0, q_1]$. That is, the seller's profit from buyers of type b is strictly increasing in q for all q between q_0 and q_1 .

Now, consider a seller with a basket price of $\hat{q} \in (q_0, q_1)$. This seller may post any price p for the cheaper good between $\hat{q} - u_b$ and u_c . In fact, for $p < \hat{q} - u_b$, the price of the more expensive good would be greater than u_b , which is never optimal. For $p > u_c$, the seller would be posting a price outside of the R_2 region, which is also never optimal. Hence, the profit of this seller is

$$S_2^*(\hat{q}) = B(\hat{q}) + \max_{p \in [\hat{q} - u_b, u_c]} C(p). \quad (46a)$$

Next, consider a seller with a basket price of q_0 . This seller may post any price p for the cheaper good between $q_0 - u_b$ and u_c . Hence, the profit of this seller is

$$S_2^*(q_0) = B(q_0) + \max_{p \in [q_0 - u_b, u_c]} C(p). \quad (47)$$

Since $S_2^*(q_0) \geq S_2^*(\hat{q})$ and $B(q_0) < B(\hat{q})$, it must be that $\max_{p \in [q_0 - u_b, u_c]} C(p)$ is strictly greater than $\max_{p \in [\hat{q} - u_b, u_c]} C(p)$. In turn, this implies that $\max_{p \in [q_0 - u_b, \hat{q} - u_b]} C(p)$ is strictly greater

than $\max_{p \in [\hat{q} - u_b, u_c]} C(p)$. Since \hat{q} was chosen arbitrarily in the interval (q_0, q_1) , $C(q_0 - u_b)$ is strictly greater than $C(p)$ for all $p \in (q_0 - u_b, u_c]$.

A seller with a basket price of $q \leq q_0$ does not find it optimal to choose a price $p \in (q_0 - u_b, q_1 - u_b]$ for the cheaper good, as $C(q_0 - u_b) > C(p)$ for all $p \in (q_0 - u_b, q_1 - u_b]$. A seller with a basket price of $q \geq q_1$ does not find it optimal to post a price $p \in (q_0 - u_b, q_1 - u_b]$ for the cheaper good, as this would imply its prices are outside the R_2 region. Finally, there are no sellers with a basket price of $q \in (q_0, q_1)$, as $G(q_0) = G(q_1)$. These observations imply that the distribution $F(p)$ has a gap between $q_0 - u_b$ and $q_1 - u_b$. In turn, this implies that the function $C(p)$ is strictly increasing in p for all $p \in [q_0 - u_b, q_1 - u_b]$, which contradicts $C(q_0 - u_b) > C(p)$ for all $p \in (q_0 - u_b, u_c]$.

Second, we show that the distribution $F(p)$ has full support over the interval $[p_\ell, u_c]$. On the way to a contradiction, suppose that $F(p)$ is such that $F'(p_0) > 0$ and $F(p_0) = F(p_1)$ for some p_0 and p_1 such that $p_\ell < p_0 < p_1 \leq u_c$. Since $F(p) = F(p_0)$ for all $p \in [p_0, p_1]$, the function $C(p)$ is strictly increasing in p for all $p \in [p_0, p_1]$. By continuity of $C(p)$, there is an interval $[p_{-1}, p_0]$ such that $C(p) < C(p_1)$ for all $p \in [p_{-1}, p_0]$. A seller with a basket price of $q \leq u_b + p_0$ does not find it optimal to post a price $p \in [p_{-1}, p_0]$ for the cheaper good, as $C(p_1) > C(p)$ for all $p \in [p_{-1}, p_0]$. A seller with a basket price of $q > u_b + p_0$ does not find it optimal to post a price $p \in [p_{-1}, p_0]$ for the cheaper good, as this would imply that its prices are outside the R_2 region. Therefore, there are no sellers posting $p \in [p_{-1}, p_0]$, which contradicts $F'(p_0) > 0$.

Finally, we show that $C(p)$ is constant for all $p \in [p_\ell, u_c]$. First, suppose that $C(p)$ is strictly increasing over some interval $(p_0, p_1) \subset [p_\ell, u_c]$. If that is the case, a seller with a basket price of $q \leq u_b + p_1$ does not find it optimal to post a price $p \in (p_0, p_1)$ for the cheaper good, as it can post the price p_1 instead and attain a higher profit. Similarly, a seller with a basket price of $q > u_b + p_1$ cannot post a price $p \in (p_0, p_1)$, as this would imply that its prices are outside of the R_2 region. Hence the distribution F has a gap between p_0 and p_1 . However, we have established that the distribution F has full support over the interval $[p_\ell, u_c]$. Therefore, $C(p)$ must be weakly decreasing for all $p \in [p_\ell, u_c]$.

Now, suppose that $C(p)$ is strictly decreasing over the interval $p \in [p_\ell, u_c]$. In this case, a seller with a basket price $q \in [q, u_b + u_c]$ chooses the lowest possible price p for the cheaper good, i.e. $u_b + p$. Hence, $F(p) = G(u_b + p)/2$ for all $p \in [p_\ell, u_c]$. Moreover $F(p)$ is such that $B(u_b + p) + C(p) = S^*$ for all $p \in [p_\ell, u_c]$. After solving this equal profit condition with respect to $F(p)$, we find that $C(p)$ is strictly increasing in p over the interval $[p_\ell, u_c]$, which contradicts the assumption that $C(p)$ is strictly decreasing. The same argument can be applied to rule out the case in which $C(p)$ is strictly decreasing over some interval

$(p_0, p_1) \subset [p_\ell, u_c]$. Therefore, $C(p)$ must be weakly increasing for all $p \in [p_\ell, u_c]$. Since $C(p)$ is both weakly decreasing and weakly increasing, it must be constant for all $p \in [p_\ell, u_c]$. In turn, this implies that $B(q)$ must be constant for all $q \in [q_\ell, u_b + u_c]$. ■

C.6 Proof of Proposition 2

We established part (i) in the main text. Here we prove part (ii). To this aim, we need to show that there exists a joint distribution H that generates the marginals G and F specified in part (i), and such that, on every point on the support of H , the profit of the seller is maximized.

We begin the analysis by identifying the region where the profit of the seller are maximized. In Lemma 1, we proved that a seller never finds it optimal to post prices (p_1, p_2) with either p_1, p_2 or both p_1 and p_2 strictly greater than u_b . It is also straightforward to show that a seller never finds it optimal to post prices (p_1, p_2) with either p_1, p_2 or both strictly smaller than p_ℓ . Therefore, we only need to check the seller's profit associated to prices $(p_1, p_2) \in [p_\ell, u_b] \times [p_\ell, u_b]$.

First, we compute the seller's profit for prices (p_1, p_2) with $p_1 + p_2 \in (u_b + u_c, q_h]$. If the seller posts prices $(p_1 + p_2)$ with $p_1 + p_2 \in [q^*, q_h]$, it attains a profit of S^* , as guaranteed by the construction of G . If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (u_b + u_c, q^*)$, it attains a profit of

$$\begin{aligned} S_1(p_1 + p_2) &= \mu_b \{ \alpha + 2(1 - \alpha)[1 - G(p_1 + p_2)] \} (p_1 + p_2) \\ &< \mu_b \{ \alpha + 2(1 - \alpha)[1 - G(q^*)] \} q^* = S^*, \end{aligned} \quad (48)$$

where the second line uses the fact that $G(p_1 + p_2) = G(q^*)$ for $p_1 + p_2 \in (u_b + u_c, q^*)$.

Second, we compute the seller's profit for prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$. In Lemma 2, we showed that the seller never finds it optimal to post prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (q_\ell, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in [p, u_c]$, it attains a profit of S^* , as guaranteed by Lemma 6. Similarly, if the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (q_\ell, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$, it attains a profit of S^* . If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, q_\ell)$, $p_1 \in [p_\ell, u_b]$ and $p_2 \in (u_c, u_b]$, it attains a profit of

$$\begin{aligned} S_2(p_1 + p_2, p_1) &= \mu_b \{ \alpha + 2(1 - \alpha) \} (p_1 + p_2) + \mu_c \{ \alpha + 2(1 - \alpha)[1 - F(p_1)] \} p_1 \\ &\leq \mu_b \{ \alpha + 2(1 - \alpha) \} q_\ell + \mu_c \{ \alpha + 2(1 - \alpha)[1 - F(p_1)] \} p_1 = S^*. \end{aligned} \quad (49)$$

The first line uses the fact that the seller trades both goods with all of its potential customers of type b . The second line uses the fact that the seller would also trade both goods with all

of its potential customers of type b at the basket price q_ℓ , and the last line uses Lemma 6. Similarly, $S_2(p_1 + p_2, p_2) \leq S^*$ for all (p_1, p_2) with $p_1 + p_2 \in (2u_c, q_\ell)$, $p_1 \in (u_c, u_b]$, $p_2 \in [\underline{p}, u_c]$.

Third, we compute the seller's profit for prices $(p_1, p_2) \in [p_\ell, u_c] \times [p_\ell, u_c]$. If the seller posts such prices, it attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \{\alpha + 2(1 - \alpha)\} (p_1 + p_2) + \sum_{i=1}^2 \mu_c \{\alpha + 2(1 - \alpha)[1 - F(p_i)]\} p_i \\ &\leq \mu_b \{\alpha + 2(1 - \alpha)\} 2u_c + \mu_c \{\alpha + 2(1 - \alpha)[1 - F(u_c)]\} 2u_c \\ &= S(u_c, u_c). \end{aligned} \tag{50}$$

The first line uses the fact that the seller trades both goods to all its potential customers of type b , it trades good 1 to $\mu_c \{\alpha + 2(1 - \alpha)[1 - F(p_1)]\}$ buyers of type c , and it trades good 2 to $\mu_c \{\alpha + 2(1 - \alpha)[1 - F(p_2)]\}$ buyers of type c . The second line uses the fact that the seller would also trade both goods to all its potential customers of type b at the prices (u_c, u_c) , and that the profit that the seller makes off of buyers of type c by posting the price $p_i \in [p_\ell, u_c]$ for good $i = \{1, 2\}$ is the same it would make by posting the price u_c instead.

If and only if $S(u_c, u_c) \leq S^*$, the highest profit that the seller can attain is S^* . Using the fact that $S_2(u_b + u_c, u_c) = S^*$, we can write the condition $S(u_c, u_c) \leq S^*$ as

$$\begin{aligned} &\mu_b \{\alpha + 2(1 - \alpha)\} 2u_c + \mu_c \{\alpha + 2(1 - \alpha)[1 - G(q^*)/2]\} 2u_c \\ &\leq \mu_b \{\alpha + 2(1 - \alpha)[1 - G(q^*)]\} (u_b + u_c) + \mu_c \{\alpha + 2(1 - \alpha)[1 - G(q^*)/2]\} u_c. \end{aligned} \tag{51}$$

After substituting out $G(q^*)$, we can write the inequality above as (30).

The functions G and F are proper distribution functions if and only if

$$\frac{\mu_c}{\mu_b} > \frac{3\alpha - 2}{(2 - \alpha)u_c/u_b} - 1, \tag{52}$$

and

$$\frac{\mu_c}{\mu_b} < \frac{\alpha(1 - u_c/u_b)}{u_c/u_b}. \tag{53a}$$

Condition (52) is necessary and sufficient for $G(q^*) > 0$, and it is condition (29). Condition (53a) is necessary and sufficient for $q^* < 2u_b$, and it holds whenever condition (30) is satisfied. If and only if (52) and (53a) are satisfied, G and F are proper distribution functions. That is, the interval $[q_\ell, u_b + u_c]$ is non-empty and, over this interval, $G(q)$ is strictly increasing in q , and such that $G(q_\ell) = 0$ and $G(u_b + u_c) = G(q^*)$, where $G(q^*) \in (0, 1)$. The interval $[q^*, q_h]$ is non-empty and, over this interval, $G(q)$ is strictly increasing in q , and such that $G(q^*) = G(u_b + u_c)$ and $G(q_h) = 1$. Similarly, the interval $[p_\ell, u_c]$ is non-empty and, over this interval, $F(p)$ is strictly increasing in p and such that $F(p_\ell) = 0$ and $F(u_c) = G(q^*)/2 \in (0, 1)$.

In the main text, we established that there exists a joint distribution H that generates the marginal F for $p \in [p_\ell, u_c]$ and the marginal G for $q \in [q_\ell, u_b + u_c]$ and that has support over the region of prices (p_1, p_2) such that $p_1 + p_2 \in [q_\ell, u_b + u_c]$, and $p_1 \in [p_\ell, u_c]$, $p_2 \in (u_c, u_b]$ or $p_1 \in (u_c, u_b]$, $p_2 \in [p_\ell, u_c]$. Over this region, the seller's profit is S^* . Moreover, we established that there exists a joint distribution H that generates the marginal G for $q \in [q^*, q_h]$ and that has support over the region of prices (p_1, p_2) such that $p_1 + p_2 \in [q^*, q_h]$ and $p_1 = p_2$. Over this region, the seller's profit is S^* . ■