Federal Reserve Bank of Philadelphia

Working Papers
Research Department

WP 16-06
Revised August 2018
February 2016
https://doi.org/10.21799/frbp.wp.2016.06

Relative Price Dispersion: Evidence and Theory

Greg Kaplan
University of Chicago and NBER

Guido Menzio
New York University and NBER

Leena Rudanko
Federal Reserve Bank of Philadelphia Research Department

Nicholas Trachter
Federal Reserve Bank of Richmond

ISSN: 1962-5361

Relative Price Dispersion: Evidence and Theory*

GREG KAPLAN
University of Chicago and NBER

GUIDO MENZIO
New York University and NBER

LEENA RUDANKO
FRB Philadelphia

NICHOLAS TRACHTER
FRB Richmond

August 2018

Abstract

Relative price dispersion is defined as persistent differences in the price that retailers set for the same good relative to the price they set for their other goods. Using a large-scale dataset on prices in the US retail market, we document that relative price dispersion accounts for about 30% of the variance of prices for the same good, in the same market, during the same week. Using a search-theoretic model of the retail market, we show that relative price dispersion can be rationalized as the equilibrium consequence of a pricing strategy used by sellers to discriminate between high-valuation buyers who need to make all of their purchases in one store, and low-valuation buyers who are able to purchase different items in different stores.

JEL Codes: L11, D40, D83, E31.
Keywords: Price Dispersion, Equilibrium Product Market Search.

*We thank the editor, John Asker, and two anonymous referees for comments and suggestions which improved the quality of the paper; Bob Hall for comments on some of our previous work which inspired us to write this paper; Jim Albrecht, George Alessandria, Fernando Alvarez, Ken Burdett, Doireann Fitzgerald, Gaetano Gaballo, Matthew Gentzkow, Erik Hurst, Boyan Jovanovic, Jose Moraga-Gonzales, Stephan Seiler, Rob Shimer, Randy Wright and participants at many seminars and conferences for their feedback. This research is based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the authors and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. 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.

1 Introduction

Using a large-scale dataset about the US retail market, we document that a large fraction of the dispersion in the price of the same good, in the same geographical area, and during the same week is caused by persistent differences in the price that different retailers set for that good relative to the price they set for other goods. We refer to this phenomenon as *Relative Price Dispersion*. We develop a theory of relative price dispersion in the context of a search-theoretic model of the retail market. According to our theory, relative price dispersion is an equilibrium consequence of a pricing strategy carried out by sellers to discriminate between high-valuation buyers who need to make all of their purchases in one location, and low-valuation buyers who are able to purchase different items in different locations.

In the first part of the paper, we document the extent of relative price dispersion in the US retail market by using the Kilts-Nielsen Retail Scanner (KNRS) dataset to measure price dispersion and decompose it into different sources. The KNRS dataset contains weekly information on prices and quantities for 1.5 million goods–defined by their Universal Product Code (UPC)–at approximately 40,000 stores across 205 Designated Market Areas (DMA)–which are geographical areas of roughly the same size as Metropolitan Statistical Areas.

We measure price dispersion as the variance of prices for the same good, in the same area and during the same week. We decompose price dispersion using a novel approach that combines some methodological aspects borrowed from the empirical literature on price dispersion (see, e.g., Sorensen (2000)) with other aspects borrowed from the empirical literature on wage inequality (see, e.g., Gottschalk and Moffitt (1994)). Following Sorensen (2000), we break down the price of a particular good at a particular store during a particular week into two parts: a store component and a store-good component. The store component is defined as the average of the (normalized) price of every good sold by the store during the week. The store-good component is defined as the difference between the price and the store component of the price. Following Gottschalk and Moffitt (1994), we consider a statistical model of prices, in which both the store component and the store-good component are described by the sum of a fixed-effect, an AR process and an MA process. We then estimate the model using the auto-covariance functions for both components of prices. The estimated model allows us to carry out a detailed decomposition of price dispersion.

We find that the variance of prices for the same good in the same geographical area and during the same week is, on average, 2.34% (equivalent to a standard deviation of 15.3%). According to the statistical model, the variance of the store component accounts for 15% of price dispersion, while the variance of the store-good component accounts for the remaining 85%. That is, price dispersion is mainly caused by dispersion in the price of a particular good relative to the price of other goods across different stores, and not by dispersion in the average price of goods across different stores. According to the statistical model, the variance of the store component of prices is almost entirely due to persistent differences across stores. In contrast, two thirds of the variance of the store-good component of prices is due to transitory differences, and one third to persistent differences. Thus, approximately 30% of price dispersion is caused by persistent differences in the relative price of the good relative to the price of other goods across different stores. This component of price dispersion is what we call relative price dispersion.

There are several theories that explain the existence of short-term differences in the price of the same good across sellers that are, on average, equally expensive. For example, according to the theory of intertemporal price discrimination (see, e.g., Conlisk, Gerstner, and Sobel (1984), Sobel (1984), Menzio and Trachter (2018)), a seller finds it optimal to occasionally lower the price of a good in order to effectively discriminate between two types of buyers: low-valuation buyers who are flexible with respect to their shopping time and, hence, can take advantage of temporary price reductions, and high-valuation buyers who cannot shop flexibly and, thus, cannot take advantage of temporary sales. If different sellers time their sales differently, there will be short-term differences in the price of the same good across sellers that are, on average, equally expensive. According to the theory of inventory management (see, e.g., Aguirregabiria (1999)), a seller finds it optimal to increase the price of a good when the inventories are low, and to lower it when the inventories are replenished. If different sellers re-stock at different times, there will be short-term differences in the price of the good across equally-expensive stores.

In contrast, there are no theories designed to explain the existence of relative price dispersion, i.e. long-term differences in the price of the same good across sellers that are, on average, equally expensive. The second part of the paper proposes a theory of relative price dispersion. The theory is cast in the context of a search-theoretic retail market in which buyers and sellers respectively demand and supply multiple (2) goods. Sellers are homogeneous. Buyers are heterogeneous with respect to their valuation for the goods and with respect to their ability to purchase different goods from different sellers. In particular, one type of buyer has a high valuation for the goods and needs to purchase all the goods in the same location. The other type of buyer has a low valuation for the goods and is able to purchase different goods at different locations. The first type of buyer describes high-wage individuals, who have a relatively low valuation for money—and, hence, value the goods relatively more—and a high valuation for time—and, hence, are unwilling to visit multiple sellers to purchase the goods. We refer to this type of buyer as busy, or type b. The second type of buyer describes low-wage individuals. We refer to this type of buyer as cool, or type

c. The retail market is subject to search frictions as in Butters (1977) and Burdett and Judd (1983). In particular, buyers have access to one randomly-selected seller with some probability, and to multiple (2) sellers otherwise.

For some parameter values, the equilibrium of the retail market features relative price dispersion, as well as dispersion in the average seller's price. Relative price dispersion is an immediate consequence of two properties of equilibrium: (i) some sellers find it optimal to set different prices for the two goods; (ii) for every seller setting a lower price for one good, there is another seller setting a lower price for the other good. The first property of equilibrium is easy to understand. Suppose that a seller sets the same price for both goods and that the common price for the two goods lies in between the valuation of type-c and type-b buyers. Clearly, the seller only trades with buyers of type b. Now, suppose that the seller decides to lower the price of first good and increase the price of the second good so as to keep the average price unchanged. At these new prices, the seller makes the same number of sales to buyers of type b, as these buyers must purchase everything in the same location and, hence, only care about the average price of the two goods. But, if the price of the first good is low enough, the seller can also trade with buyers of type c. Therefore, by setting different prices for the two goods, the seller can increase its profit. The second property of equilibrium is also easy to understand. If in equilibrium there are more sellers charging a relatively low price for one good than sellers charging a relatively low price for the other good, some sellers could increase their profits by switching the price tags of the two goods.

According to our theory, relative price dispersion is the equilibrium consequence of a strategy of price discrimination carried out by sellers. The fact that buyers of type b have a higher valuation for the goods than buyers of type c gives sellers a reason to price discriminate. The fact that buyers of type b must purchase all the goods in one location and buyers of type c can purchase different goods in different locations gives sellers an instrument to price discriminate. Indeed, sellers can effectively price discriminate between the two types of buyers by setting asymmetric prices for the two goods. In doing so, sellers can simultaneously charge a high price to the high-valuation buyers of type b who need to purchase both goods in the same location, and charge a low price (on one good) to the low-valuation buyers of type c who can purchase different goods in different locations.

We envision our model as a step towards building a comprehensive theory of price dispersion in the retail market. A comprehensive theory of price dispersion should explain why there is dispersion in the price of a particular good across stores in the same market. It should also explain why price dispersion has different sources, i.e. differences in the store component, persistent differences in the store-good component, and temporary differences in the store-good component of prices. Our model falls short of this goal because it does not explain temporary differences in the store-component of prices. However, this could be remedied by extending the model to incorporate the standard theory of temporary sales as a mean of intertemporal price discrimination. A comprehensive theory of price dispersion could then be estimated and its predictions about the shopping pattern of buyers confronted with the data (see Kaplan, Menzio, Rudanko, and Trachter (2016) for a preliminary take on this task). Once quantified and validated, the theory could be used to shed light on issues in marketing (what pricing strategies should stores adopt?), in industrial organization (welfare effect of regulations in the retail sector), and in macroeconomics (measuring correctly individual consumption as pointed out by Aguiar and Hurst (2005), measuring correctly inequality as pointed out by Attanasio and Pistaferri (2016), or understanding business cycles as pointed out by Kaplan and Menzio (2016)).

The first part of the paper is a contribution to the empirical literature on price dispersion. Our paper is the first to use a large-scale dataset that covers multiple products, each sold at multiple stores, and each observed over a long period of time. These features of the dataset allows us to borrow some methodologies from labor economics and carry out a large-scale decomposition of price dispersion into transitory and persistent differences in the store component of prices, and transitory and persistent differences in the store-good component of prices. The decomposition reveals a previously undocumented phenomenon: a large fraction of price dispersion is caused by persistent differences in the price of a particular good relative to the price of other goods across different stores.[1] Most of the previous studies on price dispersion have been focused on specific 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), 4 academic textbooks in Hong and Shum (2006), gasoline in Lewis (2006), illegal drugs in Galenianos, Pacula, and Persico (2012), mortgage brokerage

[1] A predecessor of our paper is Kaplan and Menzio (2015), which breaks down the price of a particular good in a particular 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), and (iii) a transaction component (defined residually). They then break down price dispersion into differences in the store component, differences in the store-good component and differences in the transaction component of prices. Their decomposition of price dispersion is, in spirit, similar to ours. The differences in the transaction component of prices capture, somewhat imprecisely, transitory differences in the store-good component, and differences in the store-good component capture persistent differences of the store-good component. The main problem with the decomposition in Kaplan and Menzio (2015) is that, having a much smaller dataset, they end up with very few observations for the price of a good, in a store during a quarter. As a result, their estimates of the permanent and transitory parts of the store-good component of the price are very noisy. A successor of our paper is Gorodnichenko, Sheremirov, and Talavera (2016). The focus of their paper is on the properties of online prices. Among other things, they find that relative price dispersion is an important contributor to price dispersion.

services in Woodward and Hall (2012), and many more. A few previous studies have used small-scale data covering 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). In contrast to ours, none of these papers has decomposed price dispersion into temporary and persistent differences of store and store-good components and, hence, none of these papers has documented relative price dispersion.

The second part of the paper is a contribution to search-based theories of price dispersion. We consider a version of the canonical model of Burdett and Judd (1983) in which buyers and sellers trade two goods rather than a single good, and buyers are ex-ante heterogeneous rather than homogeneous. In particular, buyers are heterogeneous with respect to their valuation for the goods and their ability to purchase different goods in different locations. We show that, for some parameter values, the equilibrium features relative price dispersion because sellers set asymmetric prices for the two goods in order to discriminate between high and low-valuation buyers. There are many search-theoretic models of price dispersion in retail markets for a single good. In Butters (1977), Varian (1980), Burdett and Judd (1983), Stahl (1989), price dispersion emerges because some buyers are in contact with a single seller and others are in contact with multiple sellers. In Reinganum (1979) and Albrecht and Axell (1984), price dispersion emerges because of heterogeneity among buyers or sellers. In Stiglitz (1987) and Menzio and Trachter (2015), price dispersion emerges because sellers are large and, hence, have an impact on the reservation price of buyers. There is a handful of models of price dispersion in retail markets for multiple goods, such as McAfee (1995), Zhou (2014), Baughman and Burdett (2015), Rhodes (2015), and Rhodes and Zhou (2015). McAfee (1995) and Baughman and Burdett (2015) also admit equilibria featuring relative price dispersion. However, in these models, buyers are homogeneous and, hence, the source of relative price dispersion in these models is not price discrimination as it is in our model. The findings in Kaplan and Menzio (2015) about significant and systematic differences in the number of stores visited and in the prices paid by different households suggests that buyers are, in fact, heterogeneous and, hence, sellers may want to discriminate them.

A strand of the Industrial Organization literature focuses on explaining why some sellers price some of their goods below cost, a phenomenon known as loss-leader. The typical explanation (see, e.g., Lal and Matutes (1994)) is cast in the context of a Hotelling (1929) model and is based on the idea that, if buyers are only aware of the price of a small subset of goods when deciding where to shop, the competition between sellers takes place only over the price of these goods. In equilibrium, the prices of the goods that are observed by buyers can be driven below cost, while the prices of the goods that are not observed by buyers remain stuck at their monopoly level. This loss-leader theory implies that sellers set the price of different goods asymmetrically. Yet, the theory does not explain why different sellers would have different loss-leader goods and, hence, why there is relative price dispersion. Lal and Matutes (1989) and Chen and Rey (2012) provide an alternative loss-leader theory which may lead to relative price dispersion. They consider the pricing strategies of duopolists facing buyers who, as in our model, differ with respect to their ability to purchase different items in different locations. They show that sellers may price goods asymmetrically and some equilibria feature relative price dispersion. While the demand side in these models is similar to the demand side in our model, the supply side and the nature of trading frictions are very different and, as a result, so is the structure of equilibrium (e.g., in our model, prices are above cost, price distributions are atomless, equilibrium features dispersion in the average price of stores and in relative prices, etc. . .).[2]

2 Evidence

In this section, we jointly analyze the dispersion of prices for the same good 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 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 geographical area. We find that a substantial fraction of the variance of the price for the same good is caused by persistent differences in the price that different stores set for that good relative to the price that they set for the other goods. We refer to this phenomenon as relative price dispersion.

2.1 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).

[2]It is also worth mentioning a Marketing literature that frames the retailer's pricing decision as a choice between Every Day Low Pricing (EDLP) and Promotional Pricing (PROMO). EDLP is a strategy that involves always selling at "low" prices. PROMO is a strategy that involves 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 with different pricing strategies. Another part of this literature (see, e.g., Ellickson and Misra (2008)) studies the retailers' choice between these two pricing strategies. This literature does not offer an explanation of 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?"

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 the log of p_{jst} as

$$\log p_{jst} = \mu_{jt} + y_{st} + z_{jst}. \tag{2.1}$$

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 z_{jst}:

$$\begin{aligned}
y_{st} &= y_s^F + y_{st}^P + y_{st}^T, & z_{jst} &= z_{js}^F + z_{jst}^P + z_{jst}^T,\\
y_{st}^P &= \rho_y y_{s,t-1}^P + \eta_{s,t}^y, & z_{jst}^P &= \rho_z z_{js,t-1}^P + \eta_{js,t}^z,\\
y_{st}^T &= \varepsilon_{s,t}^y + \sum_{i=1}^{q} \theta_{y,i}\varepsilon_{s,t-i}^y, & z_{jst}^T &= \varepsilon_{js,t}^z + \sum_{i=1}^{q} \theta_{z,i}\varepsilon_{js,t-i}^z,\\
y_s^F &= \alpha_s^y, & z_{js}^F &= \alpha_{js}^z,
\end{aligned} \tag{2.2}$$

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

[3] We 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 different goods.

[4] In 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.

Figure 1: Auto-correlation function of prices

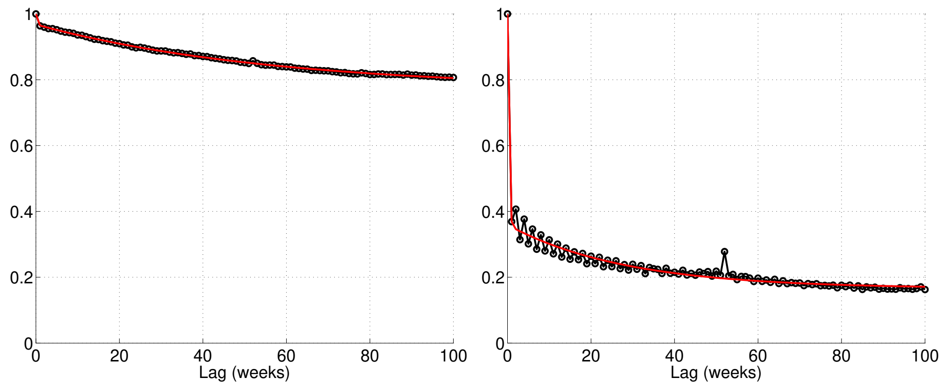

(a) Store component (b) Store-good component

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.

α_s^y and α_{js}^z are random variables with mean zero and variance $\sigma^2_{\alpha^y}$ and $\sigma^2_{\alpha^z}$. 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 $\eta^y_{s,t}$ and $\eta^z_{js,t}$ are the innovations to the AR(1) part and are assumed to be random variables with mean zero and variance $\sigma^2_{\eta^y}$ and $\sigma^2_{\eta^z}$. 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 $\varepsilon^y_{s,t}$ and $\varepsilon^z_{js,t}$ are the innovations to the MA(q) part and are assumed to be normal random variables with mean zero and variance $\sigma^2_{\varepsilon^y}$ and $\sigma^2_{\varepsilon^z}$. 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 (2.2) 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 $t = 1, ...T$, 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.,

$$\hat{\mu}_{jt} = \frac{1}{S}\sum_{s=1}^{S} \log p_{jst}. \tag{2.3}$$

We construct normalized prices as

$$\tilde{p}_{jst} = \log p_{jst} - \hat{\mu}_{jt}. \tag{2.4}$$

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.,

$$\hat{y}_{st} = \frac{1}{n_{jst}}\sum_{j=1}^{n_{jst}} \tilde{p}_{jst}, \tag{2.5}$$

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

$$\hat{z}_{jst} = \tilde{p}_{jst} - \hat{y}_{st}. \tag{2.6}$$

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.

2.2 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 6 in Appendix B shows how these 1000 products are distributed across the 10 departments defined by Nielsen. Table 7 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 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

[5] Nielsen 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.

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.

2.3 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. 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 patters in the behavior of sales that cannot be captured by our statistical model).

Table 2: Dispersion in prices: persistent and transitory

	Variance		Percent		Std. Dev.
Store component					
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
Store-good component					
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
Total		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.

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 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 the variance of prices for the same good is caused by persistent differences in the price that different stores set for that good relative to the price they set for the other goods. 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 rule out some simple explanations 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).[6] 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.[7]

3 Theory

In this section, we develop a theory of relative price dispersion. In section 3.1, we describe a model of the retail market, in which homogeneous sellers trade multiple goods to heterogeneous buyers. In section 3.2, we characterize a particular type of equilibrium. We show that, in this equilibrium, relative price dispersion emerges as the consequence of a strategy used by retailers to 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. We then find necessary and sufficient conditions for the existence of such an equilibrium.

3.1 Environment

We consider a retail market where buyers and sellers trade two products, which we refer to as good 1 and good 2. On one side of the market, there is a measure s of homogeneous sellers, with $s > 0$. Each seller can produce goods 1 and 2 at a constant marginal cost, which we normalize to zero. Each seller posts a pair of prices $(p_1, p_2) \in \mathbb{R}^2_+$, where p_1 denotes the price

[6] In principle, a wholesaler might charge different prices for different goods to different retailers. The asymmetric pricing strategy of the wholesaler could result in relative price dispersion in the retail market. Garcia and Janssen (2017) provide theoretical examples in which this pricing strategy is optimal. Unfortunately, we do not have enough data to either rule out or validate their theory.

[7] Daruich and Kozlowski (2017) carry out a decomposition of price dispersion similar to ours using retail prices from Argentina. Among other things, they find that most of price dispersion is due to differences in prices across chains of retailers rather than to differences in prices across stores of the same chain. This finding suggests that perhaps one should interpret sellers in our model as chains rather than stores.

of good 1 and p_2 denotes the price of good 2 measured in units of money.[8] Each seller chooses (p_1, p_2) so as to maximize its profit, taking as given the cumulative distribution $H(p_1, p_2)$ of prices across all sellers. We find it useful to denote with $F_i(p)$ the cumulative distribution of prices for good $i = \{1, 2\}$ across sellers and with $\lambda_i(p)$ the fraction of sellers with a price p for good $i = \{1, 2\}$. Similarly, we denote with $G(q)$ the cumulative distribution of basket prices $q = p_1 + p_2$ across sellers and with $\nu(q)$ the fraction of sellers with a basket price of q.

On the other side of the retail market, there is a measure 1 of heterogeneous buyers. There are two types of buyers, which we denote as type b for "busy" and type c for "cool." The measure of buyers of type b is μ_b, with $\mu_b \in (0, 1)$. A buyer of type b demands a single unit of good 1 and a single unit of good 2. The type-b buyer has valuation $u_b > 0$ for good 1 and good 2 and has linear utility for money. Thus, if he purchases goods 1 and 2 at the prices p_1 and p_2, the type-b buyer enjoys a utility of $2u_b - p_1 - p_2$. If he purchases only good i at the price p_i, the type-b buyer enjoys a utility of $u_b - p_i$. If he does not purchase any goods, he enjoys a utility of zero. The measure of buyers of type c is μ_c, with $\mu_c = 1 - \mu_b$. A buyer of type c demands a single unit of good 1 and a single unit of good 2. A type-c buyer has valuation u_c for good 1 and good 2 and has linear utility for money, with $u_c > 0$ and $u_c < u_b$.

Trade in the retail market is frictional, in the sense that a buyer cannot purchase the goods from any of the sellers, but only from those to which he has access. In particular, a type-b buyer has access to 1 randomly-selected seller with probability α_b, and to 2 randomly-selected sellers with probability $1 - \alpha_b$, with $\alpha_b \in (0, 1)$. We refer to the buyer as "captive" if he has access to 1 seller and as non-captive if he has access to 2 sellers. Whether he is captive or not, a type-b buyer must make all of his purchases at the same location. A type-c buyer has access to 1 randomly-selected seller with probability α_c, and to 2 randomly-selected sellers with probability $1 - \alpha_c$, with $\alpha_c \in (0, \alpha_b]$. If he is captive, a type-c buyer must make all of its purchases from that seller. However, if he is not captive, a type-c buyer can purchase one good from one seller and the other good from the other seller to which he has access.

Overall, type-b buyers differ from type-c buyers along three dimensions: (i) type-b buyers have a higher valuation for the goods; (ii) type-b buyers are more likely to have access to a single seller rather than to multiple sellers; (iii) type-b buyers must make all of their purchases from the same seller. It is natural to think of type b as high-wage buyers—who thus have a low marginal value of money and a high marginal value of time—and of type c as

[8] We assume that sellers must post positive prices. This is without loss in generality. If a buyer meets a seller with a negative price, he purchases an arbitrarily large quantity of the good and the seller's profit would be arbitrarily negative.

low-wage buyers—who thus have a high marginal value of money and a low marginal value of time. As we shall see, our theory of relative price dispersion only requires that buyers of type b and c differ along dimensions (i) and (iii). Therefore, in the main text, we assume that $\alpha_b = \alpha_c = \alpha$ in order to keep the analysis as simple as possible. In Appendix D, we consider the general case with $\alpha_b, \alpha_c \in (0, 1)$.

The model is a version of Butters (1977) and Burdett and Judd (1983) because, as in those models, we assume that there is a continuum of sellers but each buyer can only access a discrete subset of randomly-selected sellers—and the probabilities that the subset contains only one or more than one seller are both strictly positive.[9] As it is well-known, this assumption leads to price dispersion. In contrast to Butters (1977) and Burdett and Judd (1983), we assume that buyers and sellers trade two (symmetric) goods rather than a single good. We need this assumption because, in order to talk about dispersion in relative prices, we need sellers to trade at least two goods. Furthermore, in contrast to Butters (1977) and Burdett and Judd (1983), we assume that buyers are heterogeneous with respect to their valuation for the goods and their ability to purchase different items in different locations. Moreover, these two traits are correlated, in the sense that the buyers who have a higher valuation for the goods are those who are unable to purchase different items in different locations. The heterogeneity in buyers' valuations gives sellers an incentive to price-discriminate. The correlation between buyers' valuations and ability to shop at different locations gives sellers the means to price discriminate by setting a high price for one good and a low price for the other, leading in equilibrium to relative price dispersion.

The model is static, as are the models of Butters (1977) and Burdett and Judd (1983). We choose to interpret the equilibrium price dispersion in the model as a long-term outcome. That is, we choose to interpret the dispersion in the average price of sellers in the model to be the analogue of the dispersion in the persistent part of the store component of prices in the data, and the dispersion in the relative price of sellers in the model to be the analogue of the dispersion in the persistent component of the store-good component of prices in the data. We believe that this interpretation is the correct one. To see why, notice that, in a repeated version of the model, an individual seller is indifferent between keeping its prices constant and redrawing them from the equilibrium distribution every period. Therefore, in the presence of any price-adjustment cost, an individual seller strictly prefers keeping its

[9]We interpret the assumption that a buyer has only access to a discrete sample of sellers as a physical constraint (i.e., a buyer on a given day is able to reach only a few sellers) rather than an informational constraint (i.e., a buyer on a given day is aware of only a few sellers). In this sense, we interpret our model as a version of Hotelling (1929) in which buyers have a transportation cost of zero to reach some sellers, and an infinite cost to reach others.

prices constant.[10]

3.2 Discrimination Equilibrium

We restrict attention to a particular type of equilibrium, which we refer to as a Discrimination Equilibrium (henceforth DE). A DE is defined to be an equilibrium in which there is a positive measure of sellers that choose a basket price q in the top range $(u_b + u_c, 2u_b]$, a positive measure of sellers that choose a basket price q in the intermediate range $(2u_c, u_b + u_c]$, and a measure zero of sellers that choose a basket price q in the bottom range $[0, 2u_c]$. We refer to this equilibrium as a Discrimination Equilibrium because it has the feature that some sellers (those in the middle interval) want to price the two goods asymmetrically in order to discriminate between high-valuation buyers who need to purchase the whole basket of goods in the same location and low-valuation buyers who are able to purchase different goods in different locations. This seller's strategy leads to relative price dispersion. We restrict attention to a DE because it is the simplest type of equilibrium that showcases our theory of relative price dispersion.

3.2.1 General properties of equilibrium

In this subsection, we establish some basic properties of equilibrium. Importantly, these properties apply not only to a DE but to any type of equilibrium. Therefore, we will be able to appeal to these properties not only as we characterize a DE, but also when we rule out the existence of other types of equilibria.

The first lemma provides a characterization of the prices (p_1, p_2) that sellers post in equilibrium conditional on their basket price q being in the range $(u_b+u_c, 2u_b]$ and conditional on their basket price being in the range $(2u_c, u_b+u_c]$. The content of the lemma is illustrated in Figure 2.

Lemma 1: In any equilibrium:

i. A seller posts prices $(p_1, p_2) \in [0, u_b]^2$.

ii. A seller with a basket price $q \in (u_b + u_c, 2u_b]$ posts prices $(p_1, p_2) \in \mathcal{A}_1$, where

$$\mathcal{A}_1 = \left\{(p_1, p_2) \in \mathbb{R}^2_+ : p_1,\ p_2 \in (u_c, u_b]\right\}. \tag{3.1}$$

[10] In a version of the model with a discrete number of sellers and without price-adjustment costs (see, e.g., Varian 1980), there is no equilibrium in which every seller keeps its prices constant. Intuitively, after the distribution of prices is observed at the end of a period, a seller would typically have the incentive to change its price to undercut some competitors. However, if number of sellers is high (so that the gain from resetting one's prices is small) and the cost of price-adjustment is sufficiently large, sellers would find it optimal to keep their prices constant over time.

iii. A seller with a basket price $q \in (2u_c, u_b + u_c]$ posts prices $(p_1, p_2) \in \mathcal{A}_2$, where

$$\begin{aligned} \mathcal{A}_2 \quad &= \left\{(p_1, p_2) \in \mathbb{R}^2_+ : p_1 \in [0, u_c],\, p_2 \in (u_c, u_b]\right\} \\ &\cup \left\{(p_1, p_2) \in \mathbb{R}^2_+ : p_1 \in (u_c, u_b],\, p_2 \in [0, u_c]\right\} \end{aligned} \tag{3.2}$$

Proof: In Appendix C.

The first part of Lemma 1 states that a seller never posts prices higher than the valuation u_b of a type-b buyer. This property of equilibrium is intuitive. If a seller posts a price $p_i > u_b$ for good i, it sells the good neither to buyers of type b nor to buyers of type c. If the seller lowers its price for good i to u_b, it sells good i to the captive buyers of type b. Since u_b is greater than the seller's cost of production, the deviation strictly increases the profit of the seller. The second part of Lemma 1 states that any seller with a basket price $q \in (u_b + u_c, 2u_b]$ posts prices (p_1, p_2) such that p_1 and p_2 are both strictly greater than u_c and smaller than u_b. This property is an immediate implication of the first part of the lemma.

The third part of Lemma 1 states that any seller with a basket price $q \in (2u_c, u_b + u_c]$ posts prices (p_1, p_2) such that one price is smaller and the other price is greater than the valuation u_c of a type-c buyer. This property of equilibrium is central to our theory of relative price dispersion. The property implies that sellers with a basket price in the interval $(2u_c, u_b + u_c]$ always find it optimal to post a relatively high price for one good and a relatively low price for the other good. In turn, the optimality of this asymmetric pricing strategy—together with the symmetry of the equilibrium price distribution—leads to equilibrium dispersion in relative prices across sellers.

There is a simple intuition behind the third part of Lemma 1. Consider a seller with a basket price q in the interval $(2u_c, u_b + u_c]$. Suppose that a seller posts prices $p_1, p_2 \in (u_c, u_b]$ for both goods. In this case, the profit of the seller is given by

$$S(p_1, p_2) = \frac{\mu_b}{s} \left[\alpha + 2(1 - \alpha) \left(1 - G(p_1 + p_2) + \nu(p_1 + p_2)/2\right)\right] (p_1 + p_2). \tag{3.3}$$

There are $\alpha\mu_b/s$ captive buyers of type b who have access to the seller. A captive buyer of type b purchases both goods from the seller at the price $p_1 + p_2 = q$ with probability 1. There are $2(1 - \alpha)\mu_b/s$ non-captive buyers of type b who have access to the seller. A non-captive buyer of type b purchases both goods from the seller at the price q either if the other retailer with which he is in contact has a basket price greater than q, an event that occurs with probability $1 - G(q)$, or if the other retailer has a basket price of q and the buyer randomizes in favor of the seller, an event that occurs with probability $\nu(q)/2$. Since $p_1, p_2 > u_c$, the seller does not trade with any buyers of type c.

Now, suppose that the seller lowers the price of good 1 to $\hat{p}_1 = u_c$ and, simultaneously, increases the price of good 2 to $\hat{p}_2 = q - u_c \leq u_b$ so as to keep its basket price constant at

Figure 2: Pricing decision of sellers

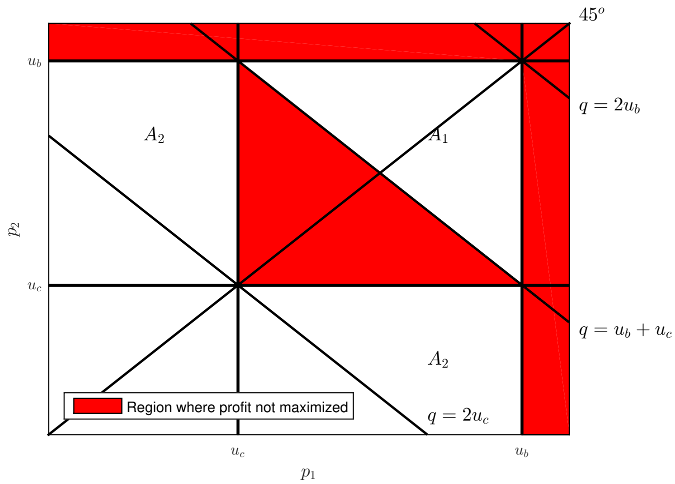

Notes: This figure considers the pricing decision of sellers discussed in the text, illustrating which regions of the (p_1, p_2) space will not be profit-maximizing. Conditional on a basket price in the interval $u_b + u_c \leq q \leq 2u_b$, sellers will not price outside $\mathcal{A}_1$. Conditional on a basket price in the interval $2u_c \leq q \leq u_b + u_c$, sellers will not price outside $\mathcal{A}_2$.

q. In this case, the profit of the seller is given by

$$\begin{aligned} S(\hat{p}_1, \hat{p}_2) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(\hat{p}_1 + \hat{p}_2) + \nu(\hat{p}_1 + \hat{p}_2)/2\right)\right](\hat{p}_1 + \hat{p}_2) \\ &+ \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)\left(1 - F_1(\hat{p}_1) + \lambda_1(\hat{p}_1)/2\right)\right]\hat{p}_1. \end{aligned} \tag{3.4}$$

The seller enjoys the same profit from trading with buyers of type b as it did before. In fact, buyers of type b make their purchasing decision based only on the price of the basket, as they have to purchase both goods in the same location, and the price of the basket is the same as before. Now, however, the seller enjoys some profit from trading with buyers of type c. There are $\alpha\mu_c/s$ captive buyers of type c who have access to the seller. A captive buyer of type c purchases good 1 from the seller with probability 1. He does not purchase good 2 because the price exceeds his valuation. There are also $2(1-\alpha)\mu_c/s$ non-captive buyers of type c who have access to the seller. A non-captive buyer of type c purchases good 1 from the seller if the other retailer with which he is in contact charges for good 1 a price greater than $\hat{p}_1$, an event that happens with probability $1 - F_1(\hat{p}_1)$, or a price equal to $\hat{p}_1$ and the buyer randomizes in favor of the seller, an event that happens with probability $\lambda_1(\hat{p}_1)/2$. Clearly, by posting the prices $(\hat{p}_1, \hat{p}_2)$, the seller enjoys a strictly higher profit than by posting (p_1, p_2).

The next lemma shows that in equilibrium there are no mass points in the distribution of basket prices, as well as in a particular region of the distribution of prices for individual goods.

Lemma 2: In any equilibrium:

i. There are no mass points in the distribution G of basket prices: $\nu(q) = 0$ for all $q \in [0, 2u_b]$;

ii. There are no mass points in the distribution F_i of prices for good i in the interval $(0, u_c]$: $\lambda_i(p) = 0$ for all $p \in (0, u_c]$.

Proof: In Appendix C.

The first part of Lemma 2 states that in equilibrium there are no mass points in the distribution of basket prices. The intuition behind this property is exactly the same as in Butters (1977) and Burdett and Judd (1983).[11] The second part of Lemma 2 states that in equilibrium there are no mass points in the distribution of prices for an individual good in the interval $(0, u_c]$. The intuition behind this property is similar to Butters (1977) and Burdett and Judd (1983). However, we cannot rule out the existence of mass points at a price $p_0 > u_c$. Intuitively, when the price of an individual good is greater than u_c, the seller only trades the good to buyers of type b who only care about basket prices. For this reason, a price $p_0 > u_c$ has no allocative role and the standard argument for ruling out mass points does not apply. We also cannot rule out the existence of a mass point at a price $p_0 = 0$. Intuitively, when sellers trade multiple goods, the fact that they set one price to zero does not imply that their profits are zero and, hence, that they could do better by resetting their price to the buyers' valuation.

3.2.2 Price distribution in the top range

We now focus on the properties that are specific to a DE. We start by characterizing the price distribution among sellers with a basket price q in the top range $(u_b + u_c, 2u_b]$.

Consider a seller with a basket price $q \in (u_b + u_c, 2u_b]$. The profit of the seller is given by

$$S(p_1, p_2) = \frac{\mu_b}{s} \left[\alpha + 2(1-\alpha)\left(1 - G(p_1 + p_2)\right)\right](p_1 + p_2)\,. \tag{3.5}$$

[11]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 strictly increase its profit. In fact, by lowering one price by an arbitrarily small amount, the seller could increase its volume of sales by $(1-\alpha)\mu_b\nu(q_0)/2 > 0$, as it could sell to all rather than to 1/2 of the non-captive buyers of type b who have access to another retailer with a basket price of q_0. Moreover, if there is a mass point at $q_0 = 0$, a seller with a basket price of q_0 could strictly increase its profit by raising both of its prices. In fact, by raising its prices from $p_1 = p_2 = 0$ to $\hat{p}_1 = \hat{p}_2 = u_b$, the seller would trade to captive buyers of type b at a positive profit rather than making no profit by selling at cost.

The above expression makes use of both Lemma 1 and Lemma 2. Lemma 1 guarantees that the seller posts prices $(p_1, p_2) \in \mathcal{A}_1$. Since $(p_1, p_2) \in \mathcal{A}_1$ implies $p_1, p_2 \leq u_b$, the seller trades both goods whenever it makes a sale to a buyer of type b. Since $(p_1, p_2) \in \mathcal{A}_1$ implies $p_1, p_2 > u_c$, the seller does not trade any of the goods to a buyer of type c. Lemma 2 guarantees that there are no mass points in the distribution G of basket prices. From these observations, it follows that the seller makes $\mu_b \left[\alpha + 2(1-\alpha)(1 - G(p_1 + p_2))\right]/s$ sales to buyers of type b and enjoys a profit of $p_1 + p_2$ on each sale, and the seller makes no sales to buyers of type c. In fact, there are $\alpha\mu_b/s$ captive buyers of type b who have access to the seller and each of them purchases both goods from the seller with probability 1. There are $2(1-\alpha)\mu_b/s$ non-captive buyers of type b who have access to the seller and each of them purchases both goods from the seller if the other retailer to which he has access has a basket price greater than $p_1 + p_2$, an event that occurs with probability $1 - G(p_1 + p_2)$. The probability that the other retailer has also a basket price of $p_1 + p_2$ is zero.

The profit of the seller in (3.5) only depends on its basket price $p_1 + p_2 = q$ and not on the individual prices p_1 and p_2 separately. Therefore, we can rewrite the seller's profit as

$$S_T(q) = \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q)\right)\right] q. \tag{3.6}$$

This property of the profit function is intuitive. A seller with a basket price $q \in (u_b + u_c, 2u_b]$ only trades with buyers of type b who care about the price of the basket of goods and not about the price of individual goods (conditional on both prices being lower than u_b, which is always the case in light of Lemma 1). And, since all relevant buyers care only about the basket price, the seller's profit only depends on the basket price. This property of the profit function implies that, for $q \in (u_b + u_c, 2u_b]$, our model in which buyers and sellers trade two goods is effectively equivalent to the standard model of Butters (1977) or Burdett and Judd (1983) in which buyers and sellers trade a single item: the basket comprising one unit of good 1 and one unit of good 2. Hence, for $q \in (u_b + u_c, 2u_b]$, the characterization of equilibrium is standard, and we can go over it at a quick pace.

Let Q_T denote the support of the distribution of basket prices G over the interval $(u_b + u_c, 2u_b]$. Let q^* and q_h denote, respectively, the infimum and the supremum of Q_T. In equilibrium, the supremum q_h must be equal to $2u_b$, the valuation of the basket for buyers of type b. To see why this is the case, suppose that $q_h < 2u_b$. If so, a seller with a basket price of q_h only trades with captive buyers of type b and attains a profit of $\alpha\mu_b q_h/s$. However, if the seller were to increase its basket price from q_h to $2u_b$, it would attain a profit of $\alpha\mu_b 2u_b/s$, which is greater than $\alpha\mu_b q_h/s$. Indeed, by raising its basket price, the seller would still only trade with captive buyers of type b but it would make a larger profit on each sale.

In equilibrium, there are no gaps in the support of the distribution G of basket prices over the interval $[q^*, q_h]$. To see why this is the case, suppose that there is a gap in G between the basket prices q_1 and q_2 with $q^* < q_1 < q_2 < q_h$. If so, a seller with a basket price of q_1 attains a profit of

$$S_T(q_1) = \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q_1)\right)\right] q_1. \tag{3.7}$$

However, if the seller were to increase its basket price from q_1 to q_2, it would attain a profit of

$$S_T(q_2) = \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q_2)\right)\right] q_2, \tag{3.8}$$

which is strictly greater than $S_T(q_1)$ as the fraction of competitors with a basket price smaller than q_1, $G(q_1)$, is the same as the fraction of competitors with a basket price smaller than q_2, $G(q_2)$.

We have thus established that Q_T is an interval $[q^*, 2u_b]$ with $u_b + u_c \leq q^* < 2u_b$. In equilibrium, a seller posting any basket price q in the interval $[q^*, 2u_b]$ must attain the maximum profit S^*. Applying this observation to $q = 2u_b$, we find that the maximum profit S^* is given by

$$S^* = \frac{\mu_b}{s}\alpha 2u_b. \tag{3.9}$$

Applying the observation for a generic $q \in [q^*, 2u_b]$, we obtain

$$\frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)(1 - G(q))\right] q = \frac{\mu_b}{s}\alpha 2u_b. \tag{3.10}$$

We can solve the above equal-profit condition with respect to the distribution G of basket prices and find that

$$G(q) = 1 - \frac{\alpha}{2(1-\alpha)}\frac{2u_b - q}{q}, \;\; \vee\, q \in [q^*, 2u_b]. \tag{3.11}$$

3.2.3 Price distribution in the intermediate range

We now characterize the price distribution among sellers with a basket price q in the intermediate range $(2u_c, u_b + u_c]$.

Consider a seller with a basket price $q \in (2u_c, u_b + u_c]$. If the seller's lowest price for an individual good is p_i, then its profit is given by

$$\begin{aligned} S(p_1, p_2) \;\; &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(p_1 + p_2)\right)\right](p_1 + p_2) \\ &+ \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)\left(1 - F_i(p_i)\right)\right] p_i. \end{aligned} \tag{3.12}$$

The above expression makes use of both Lemma 1 and Lemma 2. Lemma 1 guarantees that the seller posts prices $(p_1, p_2) \in \mathcal{A}_2$. Since $(p_1, p_2) \in \mathcal{A}_2$ implies $p_1, p_2 \leq u_b$, the seller trades both goods whenever it makes a sale to a buyer of type b. Since $(p_1, p_2) \in \mathcal{A}_2$ and $p_i = \min\{p_1, p_2\}$ implies $p_i \leq u_c$ and $p_{\neg i} > u_c$, the seller trades only good i whenever it makes a sale to a buyer of type c. Lemma 2 guarantees that there are no mass points in the basket price distribution G and in the good-i price distribution F_i at a price $p_i \in (0, u_c]$. Based on these observations, it follows that the seller makes $\mu_b \left[\alpha + 2(1-\alpha)(1 - G(p_1 + p_2))\right]$ sales to buyers of type b and enjoys a profit of $p_1 + p_2$ on each sale. Moreover, the seller enjoys a profit of $\mu_c \left[\alpha + 2(1-\alpha)(1 - F_i(p_i))\right] p_i$ from sales to buyers of type c. In fact, there are $\alpha\mu_c/s$ captive buyers of type c who have access to the seller and each of them purchases good i with probability 1. There are also $2(1-\alpha)\mu_c/s$ non-captive buyers of type c who have access to the seller and each of them purchases good i if the other retailer to whom they have access charges a price for good i greater than p_i, an event that occurs with probability $1 - F_i(p_i)$. Note that either the probability $\lambda_i(p_i)$ that the other retailer has a price of p_i is either zero or p_i is zero.

The seller's profit in (3.12) depends on both the price of the basket $q = p_1 + p_2$ and on the price p_i of the cheapest good i. Therefore, we can rewrite (3.12) as

$$S_M(q, p_i) = S_b(q) + S_{c,i}(p_i), \tag{3.13}$$

where

$$\begin{aligned} S_b(q) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q)\right)\right] q, \\ S_{c,i}(p_i) &= \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)\left(1 - F_i(p_i)\right)\right] p_i. \end{aligned}$$

The expression in (3.13) highlights the technical difficulty involved in solving for the middle part of the equilibrium price distribution. The standard strategy for solving the equilibrium price distribution involves: (i) showing that the support of the price distribution is an interval, because the distribution can have no gaps in the support; (ii) noting that everywhere on the support of the distribution the profit of the seller must be equal; (iii) solving the equal-profit condition on the support to recover the equilibrium price distribution. When, as in the case in (3.13), the profit function includes two price distributions, the standard strategy does not work. Intuitively, this is because the equilibrium may be such that sellers post prices along a one-dimensional manifold $q = f(p_i)$ and, if that is the case, using the equal-profit condition is not enough to recover both the manifold $q = f(p_i)$ and the price distributions.

In order to solve for the equilibrium price distributions G, F_1 and F_2, we follow a different approach, which consists in: (i) establishing that the two price distributions F_1 and F_2 are equal for all $p \in [0, u_c]$; (ii) establishing that the seller's profit $S_b(q)$ from trades with buyers of type b is constant for all q in the interval $[q_\ell, u_b + u_c]$, where q_ℓ is the infimum of the support of G; (iii) establishing that the seller's profit $S_{c,i}(p)$ from trades with buyers of type c is constant for all p in the interval $[p_\ell, u_c]$, where p_ℓ is the infimum of the support of F; (iv) using the equal-profit condition for $S_b(q)$ to recover the distribution of basket prices $G(q)$; (v) using the equal-profit condition for $S_c(p)$ to recover the distribution of good-i prices $F(p)$.

Lemma 3: In a Discrimination Equilibrium, $F_1(p) = F_2(p)$ for all $p \in [0, u_c]$.
Proof: In Appendix C.

Lemma 3 states that the good-1 and the good-2 price distributions are equal over the interval $[0, u_c]$. There is a simple intuition behind this property. Suppose that the two price distributions are different so that, say, $F_1(p_0) > F_2(p_0)$ for some $p_0 \in (0, u_c]$. In a discrimination equilibrium, a seller with a price of p_0 for good 1 must have a basket price of $q \in (2u_c, u_b + u_c]$ and, hence, its profit is given by $S_{M,1}(q, p_0)$. If the seller switches the prices for good 1 and good 2, it attains a profit of $S_{M,2}(q, p_0)$, which is strictly greater than $S_{M,1}(q, p_0)$ because $F_1(p_0) > F_2(p_0)$. In words, the seller can strictly increase its profit by switching around the price of the two goods because, in doing so, he can sell more to the buyers of type c while leaving its sales to buyers of type b unaltered. Therefore, in equilibrium, the distribution of prices for the two goods must be identical and, hence, we can drop the dependence of F, S_M and S_c from i.

Let Q_M denote the support of the basket price distribution G in the region $(2u_c, u_b + u_c]$, and let q_ℓ denote the infimum of Q_M. Similarly, let P_M denote the support of the individual good price distribution F in the region $[0, u_c]$, and let p_ℓ denote the infimum of P_M.

Lemma 4: In any DE:

i. $S_b(q)$ is constant for all $q \in Q_M$, where $Q_M = [q_\ell, u_b + u_c]$;

ii. $S_c(p)$ is constant for all $p \in P_M$, where $P_M = [p_\ell, u_c]$.

Proof: In Appendix C.

Lemma 4 states that the support of the basket price distribution G is the interval $[q_\ell, u_b + u_c]$ and that the seller's profit $S_b(q)$ from trades with buyers of type b is constant over this interval. Moreover, Lemma 4 states that the support of the good-i price distribution F is the interval $[p_\ell, u_c]$ and that the seller's profit $S_c(p)$ from trades with buyers of type c is constant over this interval. Lemma 4 gives us two equal-profit conditions which we can use to recover the price distributions G and F.

The proof of Lemma 4 is not standard. The details of it are in Appendix C. Here we want to provide a sketch of the argument. The first step is to show that the basket-price distribution G has full-support over $[q_\ell, u_b + u_c]$. If the support of G has a gap over some interval $(q_1, q_2) \subset [q_\ell, u_b + u_c]$, the profit function $S_b(q)$ is strictly increasing in q over that interval. If $S_b(q)$ is strictly increasing over (q_1, q_2), a seller with a basket price of q_1 makes less profit from trades with buyers of type b than if it were to set its basket price to $\hat{q} \in (q_1, q_2)$. Therefore, the seller must be able to make a higher profit from trades with buyers of type c. Since a seller with a basket price of $q \in [q_\ell, u_b + u_c]$ is constrained to post a price for its cheaper good in the interval $[q - u_b, u_c]$ in order to keep the price of its more expensive good below u_b, the seller makes a higher profit from trades with buyer of type c only if the profit function $S_c(p)$ is greater at $p = q_1 - u_b$ than at any $p \in (q_1 - u_b, q_2 - u_b)$. However, this implies that no seller finds it optimal to set a price for its cheaper good in the interval $(q_1 - u_b, q_2 - u_b)$. And if F has no mass there, the profit function $S_c(p)$ is strictly increasing between $q_1 - u_b$ and $q_2 - u_b$. This contradiction implies that G must have full-support over $[q_\ell, u_b + u_c]$. The second step of the proof is to show that the good-i price distribution F has full support over the interval $[p_\ell, u_c]$. The logic behind this step is the same as in Burdett and Judd (1983).

The third step of the proof is to show that the profit function $S_c(p)$ must be constant for all $p \in [p_\ell, u_c]$. In order to rule out the possibility that $S_c(p)$ is strictly increasing over some interval (p_1, p_2), we note that this would imply that no sellers find it optimal to post prices in that interval and, thus, F would have a gap in its support. Similarly, we rule out the possibility that $S_c(p)$ is strictly decreasing over some interval (p_1, p_2). Since $S_c(p)$ is constant for all $p \in [p_\ell, u_c]$, it must be the case that the other component of the profit function $S_b(q)$ is constant for all $q \in [q_\ell, u_b + u_c]$. This observation concludes the proof of Lemma 4.

Using Lemma 4, we can now characterize the equilibrium price distributions F and G. Since sellers make the same profit from trades with buyers of type c by posting any price $p \in [p_\ell, u_c]$, we have

$$\frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)(1 - F(p))\right] p = \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)(1 - F(u_c))\right] u_c. \tag{3.14}$$

Note that the measure of sellers with a price smaller than u_c for good i is equal to one-half of the measure of sellers with a basket price smaller than $u_b + u_c$, which is equal to the measure of sellers with a basket price smaller than q^*. That is, $F(u_c) = G(u_b + u_c)/2 = G(q^*)/2$. Using this observation to replace $F(u_c)$ on the right-hand side of (3.14) and then solving the equation for the good-i price distribution F, we obtain

$$F(p) = \frac{G(q^*)}{2} - \frac{\alpha + 2(1-\alpha)\left(1 - G(q^*)/2\right)}{2(1-\alpha)} \frac{u_c - p}{p}, \ \forall p \in [p_\ell, u_c]. \tag{3.15}$$

Note that the infimum p_ℓ of the support of the good-i price distribution is such that $F(p_\ell) =$

0. The solution of the equation $F(p_\ell) = 0$ with respect to p_ℓ is

$$p_\ell = \frac{\alpha + 2(1-\alpha)(1 - G(q^*)/2)}{2-\alpha} u_c. \tag{3.16}$$

Since sellers make the same profit from trades with buyers of type b by posting any basket price $q \in [q_\ell, u_b + u_c]$, we have

$$[\alpha + 2(1-\alpha)(1 - G(q))]\, q = [\alpha + 2(1-\alpha)(1 - G(u_b + u_c))]\, (u_b + u_c). \tag{3.17}$$

Solving the equation (3.17) with respect to the basket price distribution G and using the fact that $G(u_b + u_c) = G(q^*)$, we obtain

$$G(q) = G(q^*) - \frac{\alpha + 2(1-\alpha)(1 - G(q^*))}{2(1-\alpha)} \frac{u_b + u_c - q}{q}. \tag{3.18}$$

Note that the infimum q_ℓ of the support of the basket price distribution is such that $G(q_\ell) = 0$. The solution of the equation $G(q_\ell) = 0$ with respect to q_ℓ is

$$q_\ell = \frac{\alpha + 2(1-\alpha)(1 - G(q^*))}{2-\alpha} (u_b + u_c)\,. \tag{3.19}$$

In order to find q^*, note that, in equilibrium, a seller must be indifferent between posting the prices (u_b, u_b) and the prices (u_b, u_c). This equal-profit condition can be written as

$$\begin{aligned} \frac{\mu_b}{s}\alpha 2 u_b \quad &= \frac{\mu_b}{s}[\alpha + 2(1-\alpha)(1 - G(q^*))]\,(u_b + u_c) \\ &+ \frac{\mu_c}{s}[\alpha + 2(1-\alpha)(1 - G(q^*)/2)]\, u_c, \end{aligned} \tag{3.20}$$

where the above expression makes use of the fact that $G(u_b + u_c) = G(q^*)$ and $F(u_c) = G(q^*)/2$. Substituting out $G(q^*)$ with its equilibrium value in (3.11), we can solve the equal-profit condition (3.20) with respect to the unknown q^* and obtain

$$q^* = \frac{\mu_b \alpha 2 u_b (u_b + u_c) + \mu_c \alpha u_b u_c}{\mu_b \alpha 2 u_b - \mu_c\,(1 - \alpha/2)\, u_c}. \tag{3.21}$$

3.2.4 Necessary and sufficient conditions for equilibrium

In the previous subsection, we characterized the unique basket price distribution G and the unique good-i price distribution F that are consistent with a DE. Here, we derive a set of conditions which are necessary and sufficient for the existence of a DE.

Let us start by deriving some necessary conditions for existence of a DE. First, we derive a condition that is necessary for a seller not to find it profitable to deviate from a DE and post the prices (u_c, u_c), which are prices that would allow the seller to trade both of its goods to buyers of type c. To this aim, note that, if a seller posts prices (u_c, u_c), it attains a profit of

$$S(u_c, u_c) = \frac{\mu_b}{s}\,(2-\alpha)\,2u_c + \frac{\mu_c}{s}\,[\alpha + 2(1-\alpha)(1 - G(q^*)/2)]\,2u_c. \tag{3.22}$$

The above profit is non-greater than the equilibrium profit S^* in (3.9) if and only if

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha u_b - (2-\alpha)u_c}{u_b + (2-\alpha)u_c}\frac{u_b + u_c}{u_c}. \tag{3.23}$$

Second, we derive a condition that is necessary for q^* to belong to the interval $(u_b + u_c, 2u_b)$. To this aim, let us rewrite (3.21) as

$$\mu_b 2\alpha u_b \left(\frac{q^* - (u_b + u_c)}{q^*}\right) = \mu_c u_c \left(1 + \frac{\alpha}{2}\frac{2u_b - q^*}{q^*}\right). \tag{3.24}$$

The left-hand side of (3.24) is strictly increasing in q^* over the interval $[u_b + u_c, 2u_b]$, it takes the value 0 for $q^* = u_b + u_c$, and it takes the value $\alpha\mu_b(u_b - u_c)$ for $q^* = 2u_b$. The right-hand side of (3.24) is strictly decreasing in q^*, it takes a strictly positive value for $q^* = u_b + u_c$, and it takes the value $\mu_c u_c$ for $q^* = 2u_b$. Therefore, the q^* that solves (3.24) belongs to the interval $(u_b + u_c, 2u_b)$ if and only if

$$\frac{\mu_c}{\mu_b} < \frac{\alpha(u_b - u_c)}{u_c}. \tag{3.25}$$

Note that (3.25) is weaker than (3.23), in the sense that if (3.23) holds then (3.25) holds as well.

Third, we derive a condition that is necessary for q_ℓ to belong to the interval $(2u_c, u_b + u_c)$. From (3.19), it follows that q_ℓ is greater than $2u_c$ if and only if $G(q^*)$ is smaller than $[(2-\alpha)(u_b - u_c)]/[2(1-\alpha)(u_b + u_c)]$. This condition is always satisfied when (3.23) holds. From (3.19), it follows that q_ℓ is smaller than $u_b + u_c$ if and only if $G(q^*) > 0$. In turn, from (3.11) and (3.21), it follows that $G(q^*) > 0$ if and only if

$$\frac{\mu_c}{\mu_b} > \frac{\alpha 2u_b - (2-\alpha)(u_b + u_c)}{(2-\alpha)u_c}. \tag{3.26}$$

Conditions (3.23) and (3.26) are not only necessary for the existence of a DE. They are also sufficient. To see that this is the case, assume that (3.23) and (3.26) hold and consider a putative equilibrium in which the distribution G of sellers over basket prices is given as in (3.11) and (3.18) and q_ℓ and q^* are given by (3.19) and (3.21), in which every seller with a basket price of $q \in [q^*, 2u_b]$ posts prices $(p_1, p_2) = (q/2, q/2)$, and in which half of the sellers

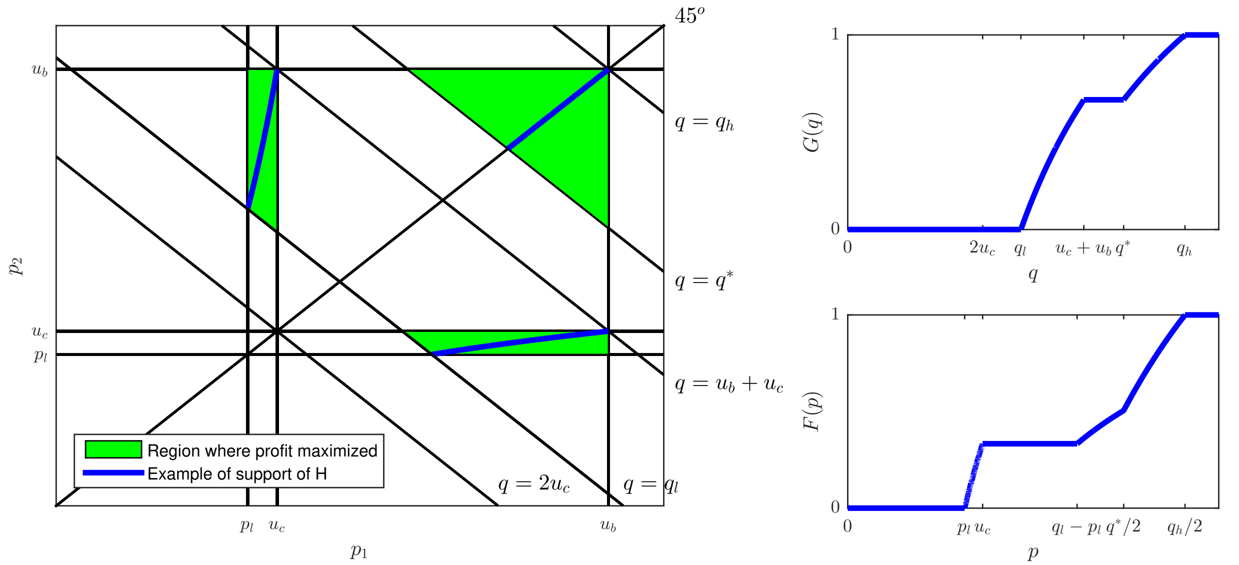

Notes: This figure shows the possible range of the support of the joint distribution $H(p_1, p_2)$, the shape of the cumulative distributions $G(q)$, and an example of the shape of the cumulative distribution $F(p)$ in the Discrimination Equilibrium.

Figure 3: Discrimination Equilibrium

with a basket price of $q \in [q_\ell, u_b + u_c]$ post prices $(p_1, p_2) = (\phi^{-1}(q), q - \phi^{-1}(q))$ and the other half posts prices $(p_1, p_2) = (q - \phi^{-1}(q), \phi^{-1}(q))$. The function $\phi(p)$ is defined as

$$\phi(p) = \frac{[\alpha + 2(1-\alpha)(1 - G(q^*))]\,(u_b + u_c)}{[\alpha + 2(1-\alpha)(1 - G(q^*))] + 2\,[\alpha + 2(1-\alpha)(1 - G(q^*)/2)]\,(u_c - p)\,/p}. \tag{3.27}$$

In words, the function $\phi(p)$ maps the lowest of the seller's prices for an individual good, p, into the seller's price for the basket of goods, $q = \phi(p)$. The function $\phi(p)$ is strictly increasing in p and such that $\phi(u_c) = u_b + u_c$. Hence, the inverse $\phi^{-1}(q)$ exists and maps the seller's basket price $q \in [q_\ell, u_b + u_c]$ into the lowest of the seller's prices for an individual good. The putative equilibrium is illustrated in Figure 3.

To see that the putative equilibrium is indeed a DE, we first need to check that G is a proper distribution function. To this aim, note that q^* belongs to the interval $(u_b + u_c, 2u_b)$ because of (3.23). Note that q_ℓ belongs to the interval $(2u_c, u_b + u_c)$, as condition (3.23) is sufficient to guarantee $q_\ell > 2u_c$ and condition (3.26) is sufficient to guarantee that $q_\ell < u_b + u_c$. Then, note that $G(q_\ell) = 0$, $G(u_b + u_c) = G(q^*)$, $G(2u_b) = 1$ and $G'(q) > 0$ for all $q \in [q_\ell, u_b + u_c] \cup [q^*, 2u_b]$. Hence, G is a proper distribution function. Since $\phi(p)$ is strictly increasing, the distribution F of prices for good i is such that $F(p) = G(\phi(p))/2$ for all $p \leq u_c$. By plugging $\phi(p)$ into G, it is immediate to verify that $F(p)$ is given by (3.15) for all $p \in [p_\ell, u_c]$.

Second, we need to check that, in the putative equilibrium, every seller with a basket price

$q \in [q^*, 2u_b]$ posts prices $(p_1, p_2) \in \mathcal{A}_1$ and that every seller with a basket price $q \in [q_\ell, u_b+u_c]$ posts prices $(p_1, p_2) \in \mathcal{A}_2$. A seller with a basket price $q \in [q^*, 2u_b]$ posts prices $(q/2, q/2) \in \mathcal{A}_1$. A seller with a basket price $q \in [q_\ell, u_b + u_c]$ posts either the prices $(\phi^{-1}(q), q - \phi^{-1}(q))$ or the prices $(q - \phi^{-1}(q), \phi^{-1}(q))$. In either case, the prices belong to $\mathcal{A}_2$, as it is easy to verify that $\phi^{-1}(q) \geq q - u_b$ and $\phi^{-1}(q) < q - u_c$.

Finally, we need to check that every seller maximizes its profit. A seller with a basket price $q \in [q^*, 2u_b]$ posts prices $(q/2, q/2) \in \mathcal{A}_1$ and, hence, its profit is given by (3.6). Moreover, by the construction of G, the seller attains a profit of $S^* = \alpha\mu_b 2u_b/s$. A seller with a basket price $q \in [q_\ell, u_b + u_c]$ posts prices $(\phi^{-1}(q), q - \phi^{-1}(q)) \in \mathcal{A}_2$ and, hence, its profit is given by (3.13). By construction of G, F and $\phi(p)$, the seller also attains a profit of S^*. It is then straightforward to verify that a seller cannot attain a profit greater than S^* by deviating from equilibrium. In particular, condition (3.23) guarantees that a seller does not find it optimal to post a basket price $q \in [0, 2u_c]$ and try to sell both goods to buyers of type c.

We have thus established the following result.

Theorem 1: A DE exists if and only if (3.23) and (3.26) hold. In any DE, G is given by (3.11) for all $q \in [q^*, 2u_b]$ and by (3.18) for all $q \in [q_\ell, u_b + u_c]$, with q_ℓ is given by (3.19) and q^* is given by (3.21). In any DE, $F_1 = F_2 = F$ and F is given by (3.15) for all $p \in [p_\ell, u_c]$, with p_ℓ given by (3.16).

A DE is unique with respect to the basket price distribution G and with respect to the good i price distribution F, as we have derived (3.11), (3.18) and (3.15) only from necessary conditions. A DE, however, is not unique with respect to the joint price distribution H. Indeed, any H that generates the marginal distributions F and G and such that every price pair on the support of H maximizes the profit of the seller is a legitimate DE. In proving the sufficiency of conditions (3.23) and (3.26), we constructed such an H by assuming that all sellers in the top section of the basket price distribution G set the same price for both goods, while sellers in the middle section of the basket distribution G set a lowest price for individual goods that is strictly increasing in their basket price. Clearly, we could have found other joint price distributions H that are consistent with equilibrium.

For any choice of H, a DE features price dispersion across sellers, in the sense that the average price of some sellers is higher than the average price of others. This follows immediately from the fact that the distribution of basket prices across sellers is non-degenerate. A DE also features relative price dispersion, in the sense that there is variation in the price of a particular good at a particular seller relative to the average price charged by that seller. This follows immediately from the fact that, in any DE, half of the sellers with a basket price $q \in (2u_c, u_b + u_c]$ have a relative price for good 1 that is greater than 1, while the other half of the sellers with a basket price $q \in (2u_c, u_b + u_c]$ have a relative price for good 1 that is smaller than 1.

The reason why a DE features price dispersion across sellers is the same as in Butters (1977) and Burdett and Judd (1983). Namely, the co-existence of captive and non-captive buyers of type b implies that a non-degenerate distribution of basket prices is required to achieve equilibrium. Now, let us explain why a DE also features relative price dispersion. Competition between sellers drives part of the distribution of basket prices to the interval $(2u_c, u_b{+}u_c]$. As established in Lemma 1, a seller with a basket price $q \in (2u_c, u_b{+}u_c]$ does not want to post the same price for both goods. Rather, the seller wants to follow an asymmetric pricing strategy, setting the price of one good below the valuation of type-c buyers and the price of the other good below the valuation of type-c buyers. As established in Lemma 3, the distribution of prices for the two goods must be symmetric and, so, if a seller posts a higher price for good 1 than for good 2, another seller must post a higher price for good 2 than for good 1. Otherwise, there would be unexploited profit opportunities. The asymmetric pricing strategy followed by each individual seller with a basket price $q \in (2u_c, u_b{+}u_c]$ together with the symmetry of the distribution of prices for the two goods guarantees the emergence of relative price dispersion.

Sellers follow an asymmetric pricing strategy to discriminate between the two types of buyers. The difference in the valuation for the goods between 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 them 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 worthwhile contrasting 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 2018). 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.

Theorem 1 states that a DE exists if and only if conditions (3.23) and (3.26) hold. Condition (3.23) is an upper bound on the measure of type-c buyers relative to the measure of type-b buyers. Intuitively, the relative measure of type-c buyers cannot be too large or else sellers would have an incentive to deviate from equilibrium and post prices that are so low that they sell both goods to buyers of type c. The upper bound is strictly positive if and only if $\alpha u_b > (2-\alpha)u_c$, meaning that condition (3.23) also requires that u_c is small enough relative to u_b. Condition (3.26) is a lower bound on the relative measure of type-c buyers. Intuitively, the measure of type-c buyers cannot be too low or else sellers would have an incentive to deviate from equilibrium and always post prices that are attractive only to buyers of type b. Notice that the lower bound in (3.26) is always smaller than the upper bound in (3.23). From these observations, we conclude that, if and only if $\alpha u_b > (2-\alpha)u_c$, there exists a population ratio $\mu_c/\mu_b > 0$ greater than (3.26) and smaller than (3.23).

Theorem 1 does not rule out the existence of other equilibria when conditions (3.23) and (3.26) hold. While we were not able to construct any other equilibrium under those conditions, one might still be concerned that relative price dispersion emerges in the DE but not in other more exotic equilibria that co-exist with the DE. To address this concern, Theorem 2 shows that, when conditions (3.23) and (3.26) hold, every possible symmetric equilibrium–i.e. every equilibrium such that $H(p_1, p_2) = H(p_2, p_1)$–features relative price dispersion.

Theorem 2: If conditions (3.23) and (3.26) hold, every symmetric equilibrium features relative price dispersion.
Proof: In Appendix C.

4 Conclusions

Using a large dataset about the US retail sector, we measured the extent and sources of dispersion in the price of the same good, in the same geographical area, during the same week. We found that about 30% of price dispersion is caused by persistent differences in the price that stores set for the same good relative to the price that they set for the other goods. We gave the name "relative price dispersion" to this source of price dispersion. After ruling out some simple explanations–such as managerial inattention, shelf management, and wholesaler-retailer relationships–we proposed a novel theory of relative price dispersion. To this aim, we considered a search-theoretic model of the retail sector, in which buyers are heterogeneous with respect to their valuation for the goods and with respect to their ability to purchase different goods at different locations. We showed that, in this retail market, relative price dispersion emerges as the equilibrium consequence of a strategy used by sellers to discriminate between high-valuation buyers who need to make all of their purchases in the same location, and low-valuation buyers who can purchase different items in different locations.

Our model is a stepping stone in the development of a comprehensive theory of price dispersion. The model can explain two of the three main sources of price dispersion: persistent differences in the average price of goods at different stores, and persistent differences in the relative price of a good relative to the price of other goods at different stores. Our model needs to be extended to a dynamic setting to explain temporary differences in the relative price of a good at different stores. The most natural way to accomplish this task would be along the lines of Menzio and Trachter (2018). Specifically, if low-valuation buyers are not only more capable of purchasing different items at different stores, but also more capable of purchasing at any time of the week, then sellers would have an incentive to run temporary sales to price discriminate between them and high-valuation buyers. Having developed a comprehensive theory of price dispersion, we could estimate it using the variance decomposition in the first part of this paper. The estimated model could then be validated by using survey data about the number of shopping trips, the number of stores visited, and the prices paid by different families.

References

Aguiar, M., and E. Hurst (2005): "Consumption versus expenditure," *Journal of Political Economy*, 113(5), 919–948.

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.

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 Monopolist," *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.

Daruich, D., and J. Kozlowski (2017): "Retail Prices: New Evidence from Argentina," Mimeo, New York University.

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.

Garcia, D., and M. Janssen (2017): "Consumer Search and Retail Channel Management," *International Journal of Industrial Organisation*, forthcoming.

Gorodnichenko, Y., V. Sheremirov, and O. Talavera (2016): "Price Setting in Online Markets: Does IT Click?," *Journal of the European Economic Association*, forthcoming.

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 (2016): "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 124(3), 771–825.

Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter (2016): "Relative Price Dispersion: Evidence and Theory," NBER Working Paper 21931.

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 Journal 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 Theory*, 67, 83–105.

Menzio, G., and N. Trachter (2015): "Equilibrium Price Dispersion with Sequential Search," *Journal of Economic Theory*, 160(6), 188–215.

——— (2018): "Equilibrium Price Dispersion Across and Within Stores," *Review of Economic Dynamics*, 28(2), 205–220.

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," *Management Science*, forthcoming.

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 3 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 3 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 that 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 4 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 3: Robustness to sample criteria

	$N_1 = 50$		$N_1 = 500$		$N_2 = 25$		$N_2 = 100$	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store								
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
Store-good								
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
Total	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, $\varepsilon^z_{js,t}$, is drawn from a skewed distribution, whose skewness we estimate alongside the other parameters of the model.[12] 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 4, 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

[12] Note 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 4: Robustness to statistical model

	MA(5)		MA(10)		Skewed MA(1)		Sales	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store-good								
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 $\delta = 52\%$.[13] The associated variance decomposition, show in Table 4, 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 ruling out some alternative 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 5, 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.

[13]As 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 9 in Appendix B, are quite different from those in the baseline sample. However, the variance decomposition for this set of goods, shown in Table 5 (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 5 (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.[14] 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 5. The low-price subsample features more relative price dispersion than the full sample: The store-good component accounts for 79%

[14]We thank Stephan Seiler for suggesting this hypothesis.

Table 5: Robustness

	Baseline		Weighted		UPC-alt		UPC-national	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store								
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
Store-good								
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/%
Store								
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
Store-good								
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/%
Store								
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
Store-good								
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 7, and 9-16 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.[15] 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 5. 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.

[15] We 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.[16] The variance decompositions for these two subsamples are shown in Table 5. 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 5, 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

[16] We 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.

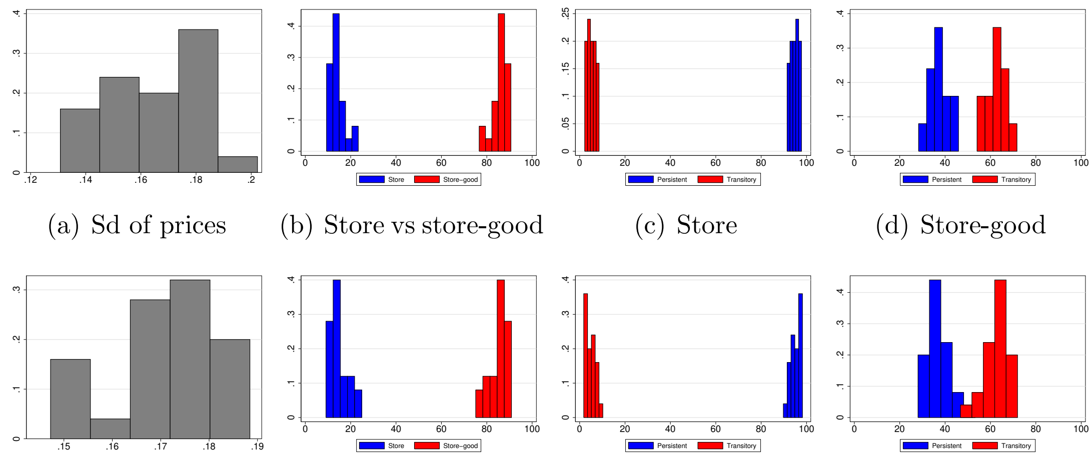

(a) Sd of prices (b) Store vs store-good (c) Store (d) Store-good

(e) Sd of prices (f) Store vs store-good (g) Store (h) Store-good

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 4: Price dispersion and variance decompositions across geographic areas

commonly purchased nationwide.

Figure 4(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 4(b), and the fraction of the variance of each component that is due to transitory versus persistent factors is shown in Figures 4(c) and 4(d), respectively. The analogous statistics are displayed for the 25 counties in the bottom row of Figure 4. Figure 4 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 6: 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 7: Share of UPCs across Product Groups: Baseline

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

Table 8: 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,

Table 9: Share of UPCs across Product Groups: Alternative

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

Table 10: Share of UPCs across Product Groups: National

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

Table 11: 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 DOEURVES-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 12: Share of UPCs across Product Groups: High Price

Product Group	Percent
CARBONATED BEVERAGES	13.65
PIZZA/SNACKS/HORS DOEURVES-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 13: 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 14: 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 15: 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 16: 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 3

C.1 Proof of Lemma 1

Proof of part (i): First, consider a seller posting prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1, p_2 > u_b$. 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, instead, the seller instead posts the prices (u_b, u_b), it sells both goods to all captive buyers of type b and it attains a profit of at least $\alpha\mu_b 2u_b/s > 0$. Therefore, a seller never finds it optimal to post prices $(p_1, p_2) \in \mathbb{R}^2_+$ such that $p_1, p_2 > u_b$.

Next, consider a seller posting prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 > u_b$, $p_2 \in (u_c, u_b]$. The seller attains a profit of

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

The above expression is easy to understand. There are $\alpha\mu_b/s$ captive buyers of type b who have access to the seller. Each one of these buyers purchases good 2 from the seller. There are also $2(1-\alpha)\mu_b/s$ non-captive buyers of type b who have access to the seller. Each one of these buyers purchases good 2 from the seller with probability $1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2$, where $\hat{G}(u_b + p_2)$ denotes the fraction of sellers that charge prices $(\tilde{p}_1, \tilde{p}_2)$ such that $\min\{\tilde{p}_1, u_b\} + \min\{\tilde{p}_2, u_b\} \leq u_b + p_2$, and $\hat{\nu}(u_b + p_2)$ is the measure of sellers that charge prices $(\tilde{p}_1, \tilde{p}_2)$ such that $\min\{\tilde{p}_1, u_b\} + \min\{\tilde{p}_2, u_b\} = u_b + p_2$. None of the buyers of type b purchases good 1 from the seller. Moreover, the seller does not trade any of the goods with buyers of type c. If, instead, the seller posts the prices (u_b, p_2), it sells both goods to its customers of type b and it attains a profit of

$$S(u_b, p_2) = \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2\right)\right](u_b + p_2). \tag{C.2}$$

Since $S(u_b, p_2) > S(p_1, p_2)$, a seller never finds it optimal to post prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 > u_b$, $p_2 \in (u_c, u_b]$. For the same reason, a seller never finds it optimal to post prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 \in (u_c, u_b]$, $p_2 > u_b$.

Finally, consider a seller posting prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 > u_b$, $p_2 \in [0, u_c]$. The seller attains a profit of

$$\begin{aligned} S(p_1, p_2) \quad &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2\right)\right] p_2 \\ &+ \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)\left(1 - F_2(p_2) + \lambda_2(p_2)/2\right)\right] p_2. \end{aligned} \tag{C.3}$$

The expression in (C.3) is the same as in (C.1) except for 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 (u_b, p_2), it attains a profit of

$$\begin{aligned} S(u_b, p_2) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2\right)\right](u_b + p_2) \\ &+ \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)\left(1 - F_2(p_2) + \lambda_2(p_2)/2\right)\right]p_2. \end{aligned} \tag{C.4}$$

Since $S(u_b, p_2) > S(p_1, p_2)$, a seller never finds it optimal to post prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 > u_b$, $p_2 \in [0, u_c]$. For the same reason, a seller never finds it optimal to post prices $(p_1, p_2) \in \mathbb{R}^2_+$ with $p_1 \in [0, u_c]$, $p_2 > u_b$. Taken together, all these observations imply that a seller always posts prices $(p_1, p_2) \in [0, u_b]^2$.

Proof of part (ii): It follows immediately from part (i).

Proof of part (iii): Suppose that there exists an equilibrium in which there is a seller that posts prices $(p_1, p_2) \in \mathbb{R}^2_+$ such that $q = p_1 + p_2 \in (2u_c, u_b + u_c]$ and $p_1, p_2 \in (u_c, u_b]$. This seller attains a profit of

$$S(p_1, p_2) = \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q) + \nu(q)/2\right)\right]q. \tag{C.5}$$

The above expression is easy to understand. There are $\alpha\mu_b/s$ captive buyers of type b who have access to the seller. Each one of these buyers purchases both goods from the seller, as $p_1 \leq u_b$ and $p_2 \leq u_b$. There are $2(1-\alpha)\mu_b/s$ non-captive buyers of type b who have access to the seller. Each one of these buyers purchases both goods from the seller with probability $1 - G(q) + \nu(q)/2$. The seller does not trade with buyers of type c because $p_1, p_2 > u_c$.

If, instead, the seller posts prices $(\hat{p}_1, \hat{p}_2) = (u_c, q - u_c)$, it attains a profit of

$$\begin{aligned} S(\hat{p}_1, \hat{p}_2) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)\left(1 - G(q) + \nu(q)/2\right)\right]q \\ &+ \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)(1 - F_1(u_c) + \lambda_1(u_c)/2)\right]u_c. \end{aligned} \tag{C.6}$$

The prices $(\hat{p}_1, \hat{p}_2)$ are such that $\hat{p}_1 + \hat{p}_2 = q$, $\hat{p}_1 \in [0, u_c]$ and $\hat{p}_2 \in (u_c, u_b]$. From these observations, it follows that the seller has the same basket price as before and, as before, both of its prices are lower than u_b. Therefore, by posting the prices $(\hat{p}_1, \hat{p}_2)$, the seller enjoys the same profit from trades with buyers of type b as it does by posting the prices (p_1, p_2). However, by posting the prices $(\hat{p}_1, \hat{p}_2)$, the seller enjoys some extra profits from trades with buyers of type c. In fact, there are $\alpha\mu_c/s$ captive buyers of type c who have access to the seller. Each one of these buyers purchases good 1 from the seller as $\hat{p}_1 \leq u_c$. There are also $2(1-\alpha)\mu_c/s$ non-captive buyers of type c who have access to the seller. Each one of these buyers purchases good 1 from the seller with probability $1 - F_1(u_c) + \lambda_1(u_c)/2$. Overall, $S(u_c, p_1 + p_2 - u_c) > S(p_1, p_2)$.

From the above observation and part (i) of the lemma, it follows that a seller with a basket price $q \in (2u_c, u_b + u_c]$ posts prices $(p_1, p_2) \in [0, u_b]^2$ such that $(p_1, p_2) \notin (u_c, u_b]^2$. This implies that the seller posts prices $(p_1, p_2) \in \mathcal{A}_2$. ■

C.2 Proof of Lemma 2

Proof of part (i): On the way to a contradiction, suppose there is an equilibrium in which the basket price distribution G has a mass point at q_0. First, let $q_0 = 0$. A seller with a basket price of q_0 must be posting prices $(p_1, p_2) = (0, 0)$ and attaining a profit of 0. If the seller deviates and posts prices $(\hat{p}_1, \hat{p}_2) = (u_b, u_b)$, it attains a profit non-smaller than $\alpha\mu_b 2u_b$. Since $\alpha\mu_b 2u_b > 0$, the seller does not maximize its profit by having a basket price of q_0.

Next, let $q_0 \in (0, 2u_b]$. A seller with a basket price of q_0 must be posting prices $(p_1, p_2) \in [0, u_b]^2$ and attaining a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \left[\alpha + 2(1-\alpha)\left(1 - G(q_0) + \nu(q_0)/2\right)\right] q_0 \\ &\quad + \textstyle\sum_{i=1}^{2} \mu_c \left[\alpha + 2(1-\alpha)\left(1 - F_i(p_i) + \lambda_i(p_i)/2\right)\right] \mathbf{1}[p_i \le u_c] p_i, \end{aligned} \tag{C.7}$$

where $\mathbf{1}[p_i \le u_c]$ is an indicator function that takes the value 1 if $p_i \in [0, u_c]$ and 0 if $p_i \in (u_c, u_b]$. If the seller deviates and posts prices $(\hat{p}_1, \hat{p}_2)$ with $0 \le \hat{p}_1 = p_1 - \epsilon_1$, $0 \le \hat{p}_2 = p_2 - \epsilon_2$, where $\epsilon_1 \ge 0$, $\epsilon_2 \ge 0$ and $\epsilon = \epsilon_1 + \epsilon_2 > 0$ arbitrarily small, it attains a profit of

$$\begin{aligned} S(\hat{p}_1, \hat{p}_2) &= \mu_b \left[\alpha + 2(1-\alpha)\left(1 - G(q_0 - \epsilon)\right)\right] (q_0 - \epsilon) \\ &\quad + \textstyle\sum_{i=1}^{2} \mu_c \left[\alpha + 2(1-\alpha)\left(1 - F_i(p_i - \epsilon_i) + \lambda_i(p_i - \epsilon_i)/2\right)\right] \mathbf{1}[p_i \le u_c](p_i - \epsilon_i). \end{aligned} \tag{C.8}$$

Clearly, $S(\hat{p}_1, \hat{p}_2) > S(p_1, p_2)$ because $(1 - G(q_0 - \epsilon)) - (1 - G(q_0) + \nu(q_0)/2)$ is greater than $\nu(q_0)/2 > 0$ and ϵ, ϵ_1 and ϵ_2 are all arbitrarily small. Therefore, the seller does not maximize its profit by having a basket price of q_0.

Proof of part (ii): On the way to a contradiction, suppose there is an equilibrium in which the good-1 price distribution F_1 has a mass point at $\tilde{p} \in (0, u_c]$. If a seller posts the price $p_1 = \tilde{p}$ for good 1, it must be posting a price $p_2 \in [0, u_b]$ for good 2. The seller attains a profit of

$$\begin{aligned} S(p_1, p_2) &= \mu_b \left[\alpha + 2(1-\alpha)[(1 - G(p_1 + p_2))\right] (p_1 + p_2) \\ &\quad + \textstyle\sum_{i=1}^{2} \mu_c \left[\alpha + 2(1-\alpha)\left(1 - F_i(p_i) + \lambda_i(p_i)/2\right)\right] \mathbf{1}[p_i \le u_c] p_i. \end{aligned} \tag{C.9}$$

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

$$\begin{aligned} S(\hat{p}_1, \hat{p}_2) &= \mu_b \left[\alpha + 2(1-\alpha)[(1 - G(p_1 + p_2 - \epsilon))\right] (p_1 + p_2 - \epsilon) \\ &\quad + \mu_c \left[\alpha + 2(1-\alpha)\left(1 - F_1(p_1 - \epsilon)\right)\right] (p_1 - \epsilon) \\ &\quad + \mu_c \left[\alpha + 2(1-\alpha)\left(1 - F_2(p_2) + \lambda_2(p_2)/2\right)\right] \mathbf{1}[p_2 \le u_c] p_i. \end{aligned} \tag{C.10}$$

Clearly, $S(\hat{p}_1, \hat{p}_2) > S(p_1, p_2)$ because $(1 - F_1(p_1 - \epsilon)) - (1 - F_1(p_1) + \lambda_1(p_1)/2)$ is greater than $\lambda_1(\tilde{p})/2 > 0$ and ϵ is arbitrarily small. Therefore, the seller does not maximzie its profit by posting a price of $\tilde{p}$ for good 1. ■

C.3 Proof of Lemma 3

As a preliminary step, we prove that, in a DE, the good-i price distribution F_i does not have a mass point over the interval $[0, u_c]$. To this aim, note that Lemma 2 guarantees that F_i does not have a mass point over $(0, u_c]$. To rule out a mass point at 0, notice that a seller with a basket price $q \in [q^*, 2u_b]$ does not set any price to zero because, in light of Lemma 1, it always posts prices $(p_1, p_2) \in (u_c, u_b]^2$. Now, consider a seller with a basket price $q \in (2u_c, u_b + u_c]$. Suppose that the seller posts prices $(p_1, p_2) = (0, q) \in \mathcal{A}_2$. In this case, the seller attains a profit of

$$S_{M,1}(q, 0) = \frac{\mu_b}{s} [\alpha + 2(1 - \alpha)(1 - G(q))] \, q.$$

The prices $(\hat{p}_1, \hat{p}_2) = (u_c, q - u_c)$ are in the region $\mathcal{A}_2$ since $\hat{p}_1 \in [0, u_c]$ and $\hat{p}_2 \in (u_c, u_b]$. Therefore, if the seller posts the prices $(\hat{p}_1, \hat{p}_2)$, it attains a profit of

$$S_{M,1}(q, u_c) = \frac{\mu_b}{s} [\alpha + 2(1 - \alpha)(1 - G(q))] \, q + \frac{\mu_c}{s} [\alpha + 2(1 - \alpha)(1 - F_1(u_c))] \, u_c.$$

Since $S_{M,1}(q, u_c) > S_{M,1}(q, 0)$, a seller with a basket price $q \in (2u_c, u_b + u_c]$ never posts any price to 0. Taken together these observations imply that, in a DE, the good-i price distribution F_i does not have a mass point over the interval $[0, u_c]$. Hence, in a DE, F_i is a continuous function over the interval $[0, u_c]$ with $F_i(0) = 0$ and $F_i(u_c) > 0$.

We now prove that $F_1 = F_2$ for all $p \in [0, u_c]$. On the way to a contradiction, suppose that $F_1(\tilde{p}) - F_2(\tilde{p}) = \delta > 0$ for some $\tilde{p} \in [0, u_c]$. Notice that $\tilde{p} > 0$ since $F_1(0) = F_2(0) = 0$. Denote as p^* the supremum of p in the interval $[0, \tilde{p}]$ for which $F_1(p^*) = F_2(p^*)$. Notice that such a p^* exists because $F_1(p) - F_2(p) = 0$ for $p = 0$, $F_1(p) - F_2(p) = \delta$ for $p = \tilde{p}$ and $F_1(p) - F_2(p)$ is continuous. Clearly, $F_1(p) - F_2(p) > 0$ over the interval $(p^*, \tilde{p})$. Moreover, a positive measure of sellers posts a price for good 1 in the interval $(p^*, \tilde{p})$ since $F_1(p^*) < F_1(\tilde{p})$.

A seller posting the price $p_1 \in (p^*, \tilde{p})$ for good 1 must be posting a price $p_2 \in (u_c, u_b]$ for good 2. The seller attains a profit of

$$\begin{aligned} S(p_1, p_2) \quad &= \frac{\mu_b}{s} [\alpha + 2(1 - \alpha)(1 - G(p_1 + p_2))] \, (p_1 + p_2) \\ &+ \frac{\mu_c}{s} [\alpha + 2(1 - \alpha)(1 - F_1(p_1))] \, p_1. \end{aligned} \tag{C.11}$$

If the seller deviates and posts the prices $(\hat{p}_1, \hat{p}_2) = (p_2, p_1)$, it attains a profit of

$$\begin{aligned} S(\hat{p}_1, \hat{p}_2) &= \frac{\mu_b}{s} \left[\alpha + 2(1-\alpha)\left(1 - G(\hat{p}_1 + \hat{p}_2)\right)\right] (\hat{p}_1 + \hat{p}_2) \\ &+ \frac{\mu_c}{s} \left[\alpha + 2(1-\alpha)\left(1 - F_2(\hat{p}_2)\right)\right] \hat{p}_2. \end{aligned} \tag{C.12}$$

Clearly, $S(\hat{p}_1, \hat{p}_2) > S(p_1, p_2)$ because $\hat{p}_1 + \hat{p}_2 = p_1 + p_2$, $\hat{p}_2 = p_1$ and $1 - F_2(\hat{p}_2) > 1 - F_1(p_1)$. Therefore, a seller does not find it optimal to post any price $p_1 \in (p^*, \tilde{p})$ for good 1, which contradicts $F_1(p^*) < F_1(\tilde{p})$. ■

C.4 Proof of Lemma 4

We first prove that the basket price distribution G has full support over the interval $[q_\ell, u_b + u_c]$. On the way to a contradiction, suppose G has a gap on its support between q_1 and q_2, with $q_\ell < q_1 < q_2 \leq u_b + u_c$. Since $G(q) = G(q_1)$ for all $q \in [q_1, q_2]$, the function $S_b(q)$ is strictly increasing in q over the interval $[q_1, q_2]$. That is, the seller's profit from trades with buyers of type b is strictly increasing in q for all $q \in [q_1, q_2]$.

Consider a seller with a basket price of q_1. This seller may post any price p for the cheaper good between $q_1 - u_b$ and u_c. In fact, for $p < q_1 - u_b$, 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 $\mathcal{A}_2$ region, which is also never optimal. Hence, the profit of this seller is

$$S_M(q_1) = S_b(q_1) + \max_{p \in [q_1 - u_b, u_c]} S_c(p). \tag{C.13}$$

Consider a seller that sets a basket price $\hat{q} \in (q_1, q_2)$. This seller may post any price p for the cheaper good between $\hat{q} - u_b$ and u_c. Hence, the maximum profit of this seller is

$$S_M(\hat{q}) = S_b(\hat{q}) + \max_{p \in [\hat{q} - u_b, u_c]} S_c(p). \tag{C.14}$$

Note that $S_M(q_1) = S^* \geq S_M(\hat{q})$ and $S_b(q_1) < S_b(\hat{q})$. Therefore, $\max_{p \in [q_1 - u_b, u_c]} S_c(p)$ must be strictly greater than $\max_{p \in [\hat{q} - u_b, u_c]} S_c(p)$. In turn, this implies that $\max_{p \in [q_1 - u_b, \hat{q} - u_b]} S_c(p)$ must be strictly greater than $\max_{p \in [\hat{q} - u_b, u_c]} S_c(p)$. Since $\hat{q}$ was chosen arbitrarily in the interval (q_1, q_2), $S_c(q_1 - u_b) > S_c(p)$ for all $p \in (q_1 - u_b, u_c]$.

Now, notice that a seller with a basket price of $q \leq q_1$ does not find it optimal to choose a price $p \in (q_1 - u_b, q_2 - u_b)$ for the cheaper good, as $S_c(q_1 - u_b) > S_c(p)$ for all $p \in (q_1 - u_b, q_2 - u_b)$. A seller with a basket price of $q \geq q_2$ does not find it optimal to post a price $p \in (q_1 - u_b, q_2 - u_b)$ for the cheaper good, as this would imply its prices are not in $\mathcal{A}_2$. Finally, there are no sellers with a basket price of $q \in (q_1, q_2)$, as $G(q_1) = G(q_2)$. These observations imply that the distribution $F(p)$ has a gap between $q_1 - u_b$ and $q_2 - u_b$. In turn, this implies that the function $S_c(p)$ is strictly increasing in p for all $p \in [q_1 - u_b, q_2 - u_b]$, which contradicts $S_c(q_0 - u_b) > S_c(p)$ for all $p \in (q_1 - u_b, u_c]$.

Second, we show that the good-i price distribution F has full support over the interval $[p_\ell, u_c]$. On the way to a contradiction, suppose that F has a gap in its support between p_1 and p_2, with $p_\ell < p_1 < p_2 \leq u_c$. Since $F(p) = F(p_0)$ for all $p \in [p_1, p_2]$, the function $S_c(p)$ is strictly increasing in p for all $p \in [p_1, p_2]$. By continuity of $S_c(p)$, there must be an interval (p_0, p_1) such that $S_c(p) < S_c(p_2)$ for all $p \in (p_0, p_1)$. Since (p_0, p_1) intersects with the support of F, there is a positive measure of sellers that post prices in (p_0, p_1). A seller with a basket prices 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 $S_c(p_2) > S_c(p)$ for all $p \in (p_0, p_1)$. A seller with a basket price of $q > u_b + p_1$ does not find it optimal to post a price $p \in (p_0, p_1)$ for the cheaper good, as this would imply that its prices are outside the $\mathcal{A}_2$ region. Therefore, there are no sellers posting a price p in the interval (p_1, p_2). We have thus reached a contradiction.

Finally, we show that $S_c(p)$ is constant for all $p \in [p_\ell, u_c]$. First, suppose that $S_c(p)$ is strictly increasing over some interval $(p_1, p_2) \subset [p_\ell, u_c]$. If that is the case, a seller with a basket price of $q \leq u_b + p_2$ does not find it optimal to post a price $p \in (p_1, p_2)$ for the cheaper good, as it can post the price p_2 instead and attain a higher profit. Similarly, a seller with a basket price of $q > u_b + p_2$ cannot post a price $p \in (p_1, p_2)$, as this would imply that its prices are outside of the $\mathcal{A}_2$ region. Hence the distribution F has a gap between p_1 and p_2. However, we have established that the distribution F has full support over the interval $[p_\ell, u_c]$. Therefore, $S_c(p)$ must be weakly decreasing for all $p \in [p_\ell, u_c]$.

Now, suppose that $S_c(p)$ is strictly decreasing over the interval $[p_\ell, u_c]$. 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 $S_b(u_b + p) + S_c(p) = S^*$ for all $p \in [p_\ell, u_c]$. After solving this equal profit condition with respect to $F(p)$, we find that $S_c(p)$ is strictly increasing in p over the interval $[p_\ell, u_c]$, which contradicts the assumption that $S_c(p)$ is strictly decreasing. The same argument can be applied to rule out the case in which $S_c(p)$ is strictly decreasing over some interval $(p_1, p_2) \subset [p_\ell, u_c]$. Therefore, $S_c(p)$ must be weakly increasing for all $p \in [p_\ell, u_c]$. Since $S_c(p)$ is both weakly decreasing and weakly increasing, it must be constant for all $p \in [p_\ell, u_c]$. In turn, this implies that $S_b(q)$ must be constant for all $q \in [q_\ell, u_b + u_c]$. ■

C.5 Proof of Theorem 2

We want to prove that every symmetric equilibrium features relative price dispersion whenever the conditions (3.23) and (3.26) hold. To this aim, let us first suppose that there exists a symmetric equilibrium without relative price dispersion in which every seller has a basket price q in the top region $(u_b+u_c, 2u_b]$. Since the equilibrium has no relative price dispersion, every seller has the same ratio ρ between the price of good 1 and the price of good 2. Since the equilibrium is symmetric, $\rho = 1$. Hence, every seller with a basket price of q posts prices $(p_1, p_2) = (q/2, q/2)$. Following the same argument as in the proof of Lemma 2, it is easy to show that the basket price distribution G must be atomless and the support of G must be an interval $[q_\ell, 2u_b]$, with $q_\ell \in (u_b + u_c, 2u_b)$. Hence, the profit for a seller posting any equilibrium prices is

$$S^* = \frac{\mu_b}{s}\alpha 2u_b. \tag{C.15}$$

If a seller deviates from the equilibrium and posts the prices (u_b, u_c), it attains a profit of

$$S(u_b, u_c) = \frac{\mu_b}{s}(2-\alpha)(u_b+u_c) + \frac{\mu_c}{s}(2-\alpha)u_c. \tag{C.16}$$

The existence of equilibrium requires $S^* \geq S(u_b, u_c)$ or, equivalently,

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha 2u_b - (2-\alpha)(u_b+u_c)}{(2-\alpha)u_c}. \tag{C.17}$$

Since (C.17) contradicts condition (3.26), there cannot exist a symmetric equilibrium without relative price dispersion such that every seller has a basket price $q \in (u_b + u_c, 2u_b]$. In fact, it is easy to show that there cannot be any equilibrium such that every seller has a basket price $q \in (u_b + u_c, 2u_b]$.

Next, suppose that there exists a symmetric equilibrium without relative price dispersion in which every seller has a basket price q in the bottom region $[0, 2uc]$. Since the equilibrium features no relative price dispersion and is symmetric, every seller with a basket price of q must post prices $(p_1, p_2) = (q/2, q/2)$. Given this observation, it is easy to prove that the basket price distribution G must be atomless and the support of G must be some interval $[q_\ell, 2u_c]$. Hence, the profit for a seller posting any equilibrium prices is

$$S^* = \frac{\mu_b + \mu_c}{s}\alpha 2u_c. \tag{C.18}$$

If a seller deviates from equilibrium and posts the prices (u_b, u_c), it attains a profit of

$$S(u_b, u_c) = \frac{\mu_b}{s}\alpha(u_b+u_c) + \frac{\mu_c}{s}\alpha u_c. \tag{C.19}$$

The existence of equilibrium requires $S^* \geq S(u_b, u_c)$ or, equivalently,

$$\frac{\mu_c}{\mu_b} \leq \frac{u_b - u_c}{u_c}. \tag{C.20}$$

Since (C.20) is violated when (3.23) holds, there cannot exist a symmetric equilibrium without relative price dispersion such that every seller has a basket price $q \in [0, 2u_c]$.

Third, suppose that there exists a symmetric equilibrium without relative price dispesion in which a positive measure of sellers have a basket price $q \in [0, 2uc]$, a positive measure of sellers have a basket price $q \in (u_b+u_c, 2u_b]$, and no sellers have a basket price $q \in (2u_c, u_b+u_c]$. Since the equilibrium features no relative price dispersion and is symmetric, every seller with a basket price of q must post prices $(p_1, p_2) = (q/2, q/2)$. Given this observation, we can show that G is atomless and its support is the union of the intervals $[q_\ell, 2u_c]$ and $[q^*, 2u_b]$. Hence, the profit for a seller posting any equilibrium prices is

$$\begin{aligned} S^* &= \frac{\mu_b}{s}\alpha 2u_b \\ &= \frac{\mu_b+\mu_c}{s}\left[\alpha + 2(1-\alpha)(1-G(q^*))\right]2u_c. \end{aligned} \tag{C.21}$$

If a seller deviates from the equilibrium and posts the prices (u_b, u_c), it attains a profit of

$$\begin{aligned} S(u_b, u_c) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)(1-G(q^*))\right](u_b+u_c) \\ &+ \frac{\mu_c}{s}(\left[\alpha + 2(1-\alpha)(1-G(q^*))\right]u_c. \end{aligned} \tag{C.22}$$

The existence of equilibrium requires $S^* \geq S(u_b, u_c)$. A little bit of algebra reveals that this condition is never satisfied. Hence, there cannot exist a symmetric equilibrium without relative price dispesion in which a positive measure of sellers have a basket price $q \in [0, 2u_c]$, a positive measure of sellers have a basket price $q \in (u_b + u_c, 2u_b]$, and no sellers have a basket price $q \in (2u_c, u_b + u_c]$.

Finally, notice that, in any symmetric equilibrium in which a positive measure of sellers have a basket price q in the intermediate region $(2u_c, u_b + u_c]$, there must be relative price dispersion because of part (iii) of Lemma 1. Hence, there cannot exist a symmetric equilibrium without relative price dispesion in which a positive measure of sellers have a basket price $q \in (2u_c, u_b + u_c]$. ■

D Discrimination Equilibrium with $\alpha_b \neq \alpha_c$

In Section 3, we established that a Discrimination Equilibrium exist when conditions (3.23) and (3.26) hold and we showed that a DE always features relative price dispersion. The analysis in Section 3 is carried out under the assumption that $\alpha_b = \alpha_c = \alpha \in (0, 1)$. In this Appendix, we want to show that a DE also exists and also features relative price dispersion under the weaker assumption that $\alpha_b, \alpha_c \in (0, 1)$. That is, we want to generalize some of the results of Section 3 to the case in which buyers of type b and buyers of type c differ not only with respect to their valuation for the goods and with respect to their ability to purchase different goods in different locations, but also with respect to their probability of being captive.

For the sake of brevity, we only derive sufficient conditions for the existence of a DE. To this aim, we are going to construct a putative DE as in Section 3 and derive conditions under which this is indeed an equilibrium. For $q \in [q^*, 2u_b]$, let the distribution G of sellers over basket prices be given by

$$G(q) = 1 - \frac{\alpha_b}{2(1-\alpha_b)} \frac{2u_b - q}{q}. \tag{D.1}$$

For $q \in [q_\ell, u_b + u_c]$, let the distribution G of sellers over basket prices be given by

$$G(q) = G(q^*) - \frac{\alpha_b + 2(1-\alpha_b)(1-G(q^*))}{2(1-\alpha_b)} \frac{u_b + u_c - q}{q}. \tag{D.2}$$

Let cutoffs q^* and q_ℓ be respectively given by

$$\begin{aligned} q^* = \quad & \left[\mu_b \alpha_b 2u_b - \mu_c \left(1 - \frac{(1-\alpha_c)\alpha_b}{2(1-\alpha_b)}\right) u_c\right]^{-1} \\ & \cdot \left[\mu_b \alpha_b 2u_b(u_b + u_c) + \mu_c \frac{(1-\alpha_c)\alpha_b}{2(1-\alpha_b)} 2u_b u_c\right], \end{aligned} \tag{D.3}$$

and

$$q_\ell = \frac{\alpha_b + 2(1-\alpha_b)(1-G(q^*))}{2-\alpha_b}(u_b + u_c) \tag{D.4}$$

Assume that every seller with a basket price of $q \in [q^*, 2u_b]$ posts prices $(p_1, p_2) = (q/2, q/2)$. Moreover, assume that half of the sellers with a basket price of $q \in [q_\ell, u_b + u_c]$ post prices $(p_1, p_2) = (\phi^{-1}(q), q - \phi^{-1}(q))$ and the other half posts prices $(p_1, p_2) = (q - \phi^{-1}(q), \phi^{-1}(q))$. The function $\phi(p)$ is given by

$$\begin{aligned} & \phi(p) \\ = \quad & \left[\frac{\alpha_c + 2(1-\alpha_c)(1-G(q^*)/2)}{1-\alpha_c} \frac{u_c - p}{p} + \frac{\alpha_b + 2(1-\alpha_b)(1-G(q^*))}{2(1-\alpha_b)}\right]^{-1} \\ & \cdot \frac{\alpha_b + 2(1-\alpha_b)(1-G(q^*))}{2(1-\alpha_b)}(u_b + u_c). \end{aligned} \tag{D.5}$$

The putative equilibrium is constructed as in Section 3. In particular, the first part of the G distribution is constructed so as to keep the seller's profit constant for all prices $(p_1, p_2) \in A_1$ such that $p_1 + p_2 \in [q^*, 2u_b]$. The second part of the G distribution is constructed so as to keep the seller's profit from trades with buyers of type b constant for all prices $(p_1, p_2) \in A_2$ such that $p_1 + p_2 \in [q_\ell, u_b + u_c]$. The function $\phi(p)$ is constructed so as to generate an F distribution that keeps the seller's profit from trades with buyers of type c

constant over the region $(p_1, p_2) \in A_2$ such that $p_1 + p_2 \in [q_\ell, u_b + u_c]$. In particular, for all $p \in [p_\ell, u_c]$, the function $\phi(p)$ generates a distribution F of prices for good-i that is given by

$$F(p) = \frac{G(q^*)}{2} - \left(\frac{\alpha_c + 2(1-\alpha_c)(1 - G(q^*)/2)}{2(1-\alpha_c)} \right) \frac{u_c - p}{p}, \tag{D.6}$$

where the cutoff p_ℓ is

$$p_\ell = \frac{\alpha_c + 2(1-\alpha_c)(1 - G(q^*)/2)}{2 - \alpha_c} u_c. \tag{D.7}$$

In order to verify that the putative DE is indeed an equilibrium, we first need to check that G is a proper distribution function. To this aim, note that the cutoff q^* belongs to the interval $(u_b + u_c, 2u_b)$ if and only if

$$\frac{\mu_c}{\mu_b} < \frac{\alpha_b(u_b - u_c)}{u_c}. \tag{D.8}$$

The cutoff q_ℓ belongs to the interval $(2u_c, u_b + u_c)$ if and only if

$$G(q^*) > 0 \iff \frac{\mu_c}{\mu_b} > \frac{\alpha_b 2u_b - (2-\alpha_b)(u_b + u_c)}{(2-\alpha_c)u_c}, \tag{D.9}$$

and

$$G(q^*) < \frac{(2-\alpha_b)(u_b - u_c)}{2(1-\alpha_b)(u_b + u_c)}. \tag{D.10}$$

Under these conditions, the G distribution is a proper distribution function. In fact, $G(q_\ell) = 0$, $G'(q) > 0$ for all $q \in [q_\ell, u_b + u_c]$, $G(u_b + u_c) = G(q^*)$, $G'(q) > 0$ for all $q \in [q^*, 2u_b]$ and $G(2u_b) = 1$.

Second, we need to verify that every seller with a basket price $q \in [q^*, 2u_b]$ posts prices $(p_1, p_2) \in \mathcal{A}_1$ and that every seller with a basket price $q \in [q_\ell, u_b + u_c]$ posts prices $(p_1, p_2) \in \mathcal{A}_2$. A seller with a basket price $q \in [q^*, 2u_b]$ posts prices $(q/2, q/2) \in \mathcal{A}_1$. A seller with a basket price $q \in [q_\ell, u_b + u_c]$ posts either prices $(\phi^{-1}(q), q - \phi^{-1}(q))$ or $(q - \phi^{-1}(q), \phi^{-1}(q))$. These prices belong to $\mathcal{A}_2$ if and only if $\phi^{-1}(q) < q - u_c$ and $\phi^{-1}(q) \geq q - u_b$ for all $q \in [q_\ell, u_b + u_c]$. The condition $\phi^{-1}(q) < q - u_c$ is always satisfied. The condition $\phi^{-1}(q) \geq q - u_b$ is satisfied if and only if

$$G(q^*) \leq \frac{2-\alpha_c}{1-\alpha_c} \frac{u_b + u_c}{u_b} - \frac{2-\alpha_b}{2(1-\alpha_b)} \frac{u_c}{u_b}. \tag{D.11}$$

Finally, we need to check that every seller maximizes its profit. A seller with a basket price $q \in [q^*, 2u_b]$ posts prices $(p_1, p_2) \in \mathcal{A}_1$ and, hence, its profit is given by

$$S_T(q) = \frac{\mu_b}{s} \left[\alpha + 2(1-\alpha)(1 - G(q)) \right] q. \tag{D.12}$$

By construction of the basket price distribution G, the seller's profit is

$$S_T(q) = \frac{\mu_b}{s}\alpha_b 2u_b \equiv S^*. \tag{D.13}$$

A seller with a basket price $q \in [q_\ell, u_b + u_c]$ posts prices $(p_1, p_2) \in \mathcal{A}_2$ and, hence, its profit is given by

$$S_M(q,p) = S_b(q) + S_c(p), \tag{D.14}$$

where

$$\begin{aligned} S_b(q) &= \frac{\mu_b}{s}\left[\alpha + 2(1-\alpha)(1-G(q))\right] q, \\ S_c(p) &= \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)(1-F(p))\right] p. \end{aligned}$$

By construction of the basket price distribution G, of the good-i price distribution F and of the function $\phi(p)$, the seller's profit $S_M(q,p)$ is equal to S^*.

A seller attains the profit S^* by posting any of the equilibrium prices. Can the seller attain a higher profit by posting off-equilibrium prices? The most profitable deviation would be to post the prices (u_c, u_c), which would give the seller a profit of

$$S(u_c, u_c) = \frac{\mu_b}{s}(2-\alpha_b)2u_c + \frac{\mu_c}{s}\left[\alpha + 2(1-\alpha)(1 - G(q^*)/2)\right] 2u_c. \tag{D.15}$$

The profit $S(u_c, u_c)$ is smaller than the equilibrium profit S^* if and only if

$$G(q^*) \leq \frac{u_b}{u_b + u_c} \Leftrightarrow \frac{\mu_c}{\mu_b} \leq \frac{\alpha_b u_b - (2-\alpha_b)u_c}{u_b + (2-\alpha_c)u_c}\frac{u_b + u_c}{u_c}. \tag{D.16}$$

Notice that condition (D.16) implies conditions (D.8) and (D.10).

Overall, the putative Discrimination Equilibrium is an equilibrium if and only if conditions (D.9), (D.11) and (D.16) hold. Condition (D.16) generalizes condition (3.23) to the case of $\alpha_b, \alpha_c \in (0,1)$ and gives an upper bound on the measure of buyers of type c relative to the measure of buyers of type b. Intuitively, the relative measure of type-c buyers cannot be too large or else sellers would have an incentive to deviate from equilibrium and post prices that are so low that they sell both goods to buyers of type c. The upper bound is strictly positive if and only if $\alpha_b u_b > (2-\alpha_b)u_c$, meaning that condition (D.16) also requires that u_c is small enough relative to u_b. Condition (D.9) generalizes condition (3.26) and gives a lower bound on the relative measure of buyers of type c. Intuitively, the measure of type-c buyers cannot be too low or else sellers would have an incentive to deviate from equilibrium and always post prices that are attractive only to buyers of type b. The lower bound in (D.9) is always smaller than the upper bound in (D.16). From these observations, it follows that, if and only if $\alpha_b u_b > (2-\alpha_b)u_c$, there is a population ratio $\mu_c/\mu_b > 0$ that satisfies both (D.9) and (D.16).

Condition (D.11) is an upper bound on $G(q^*)$. The condition guarantees that there exists a pricing schedule $\phi(p)$ that generates the marginal distributions F and G and that is consistent with the seller's profit maximization. The sufficient conditions for existence of a DE in Section 3 did not have an analogue to (D.11). This is because, when $\alpha_b = \alpha_c = \alpha$, the upper bound on $G(q^*)$ implied by (D.11) is greater than 1 and, hence, always satisfied. The same is true when $\alpha_b < \alpha_c$. Therefore, for all $\alpha_b, \alpha_c \in (0,1)$ with $\alpha_b < \alpha_c$, (D.11) holds and there exists a population ratio $\mu_c/\mu_b > 0$ that satisfies both (D.9) and (D.16) if $\alpha_b u_b > (2-\alpha_b)u_c$. When $\alpha_c < \alpha_b$, the upper bound on $G(q^*)$ implied by (D.11) is not always greater than 1 and the condition needs to be checked. However, it is easy to see that if u_c is small enough, the upper bound on $G(q^*)$ is greater than 1. Therefore, for all $\alpha_b, \alpha_c \in (0,1)$ with $\alpha_c < \alpha_b$, (D.11) holds if u_c is small enough relative to u_b and there exists a population ratio $\mu_c/\mu_b > 0$ that satisfies both (D.9) and (D.16) if $\alpha_b u_b > (2-\alpha_b)u_c$. Overall, for any $\alpha_b, \alpha_c \in (0,1)$ with $\alpha_b < \alpha_c$, there are values for the other parameters such that the sufficient conditions (D.11), (D.9) and (D.16) are satisfied and a DE exists.

Clearly, the DE that we constructed above features relative price dispersion. It is important to notice that this property is not an artifact of our particular selection for the joint price distribution H. In fact, Lemma 1 and Lemma 3 hold for generic $\alpha_b, \alpha_c \in (0,1)$. And from Lemma 1 and Lemma 3, it follows that any DE must feature relative price dispersion.
■