Relative Price Dispersion: Evidence and Theory*

Greg Kaplan
Princeton University and NBER

Guido Menzio
University of Pennsylvania and NBER

Leena Rudanko
FRB Philadelphia

Nicholas Trachter
FRB Richmond

July 2015

Abstract

We use a large dataset on retail pricing to document that a sizeable portion of the cross-sectional variation in the price at which the same good trades in the same period and in the same market is due to the fact that stores that are, on average, equally expensive set persistently different prices for the same good. We refer to this phenomenon as relative price dispersion. We argue that relative price dispersion stems from the sellers' attempt to discriminate between high-valuation buyers who need to make all of their purchases in the same store, and low-valuation buyers who are willing to purchase different items from different stores. Using a dataset on the shopping behavior of households, we provide some evidence supporting this theory.

JEL Codes: L11, D40, D83, E31.

Keywords: Price Dispersion, Equilibrium Product Market Search.

*Kaplan: Department of Economics, Princeton University, Fisher Hall, Princeton, NJ 08544 (email: gkaplan@princeton.edu); Menzio: Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19013 (email: gmenzio@sas.upenn.edu); Rudanko: Federal Reserve Bank of Philadelphia, Ten Independene Mall, Philadelphia, PA 19106 (email: leena.rudanko@gmail.com);. Trachter: Federal Reserve Bank of Richmond, 701 E. Byrd Street, Richmond, VA 23219 (email: nicholas.trachter@rich.frb.org). We are grateful to our audiences at Wharton/UPenn, the Federal Reserve Bank of Minneapolis, George Washington University, the Search and Matching workshop in Philadelphia, and the Rome Junior Conference on Macroeconomics for comments. We are grateful to the Kilts Center for Marketing and the Nielsen Company for sharing their data. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Philadelphia, Federal Reserve Bank of Richmond or the Federal Reserve System.

1 Introduction

Using a large dataset on retail pricing, we document that a significant fraction of the cross-sectional variation in the price at which the same good is sold in the same period of time and in the same geographical area is due to the fact that stores that, on average, are equally expensive set persistently different prices for that good. We refer to this phenomenon as relative price dispersion. We propose a theory of relative price dispersion in a search-theoretic model of the retail market, where buyers and sellers, respectively, demand and supply multiple goods. We argue that relative price dispersion stems from the sellers' incentive to discriminate between high-valuation buyers who need to make all of their purchases in the same store, and low-valuation buyers who are willing to purchase different items from different stores.

In the first part of the paper, we carry out a novel decomposition of the sources of dispersion in the price at which the very same good is sold in the same week and in the same geographical area. The main finding of the decomposition is that a sizeable fraction of price dispersion is due to the fact that, among stores that are equally expensive on average, the same good is sold at prices that are persistently different. That is, a sizeable fraction of price dispersion is due to persistent differences across stores in the price of the good relative to the average price of the store.

We measure and decompose the sources of price dispersion using the Kielts-Nielsen Retail Scanner Dataset (KNRS), which provides weekly price and quantity information for 1.5 million UPCs at around 40,000 stores in over 2,500 counties across 205 Designated Market Areas, which are geographical areas of roughly the same size as Metropolitan Statistical Areas. We begin by computing the average price of each good (as defined by its UPC) in a particular week and in a particular geographical area. We normalize the price at which the good is traded at a particular store by taking its (log) difference with respect to the average price in the relevant time period and geographical area. Then, we break down the normalized price into a store component—defined as the average of the normalized price of all the goods sold by the store in the relevant week—and a store-good component—defined residually as the (log) difference between the price of the good and average price of the store. The average standard deviation of normalized prices for the same good in the same time period and in the same geographical area is 15.3%.[1] Moreover, we find that only 15% of the variance of

[1]The extent of price dispersion we document using the KNRS is quite similar to what previously documented by Stigler (1961), Pratt, Wise, and Zeckhauser (1979) or Sorensen (2000) for narrow sets of goods, and by Eden (2013) or Kaplan and Menzio (2014a) using other datasets that also cover a wide variety of products.

prices is due to the variance of the store component, i.e. due to the fact that the same good is sold at stores that have different average prices, while 85% is due to the variance of the store-good component, i.e. due to the fact that the same good is sold at different prices at stores that are equally expensive.

We then break down the store and the store-good component of prices into temporary and persistent parts. To this aim, we follow an approach that has been commonly applied in the literature on labor economics to decompose wage inequality (see, e.g, Gottschalk and Moffitt (1994) and Blundell and Preston (1998) but that had never been used to study price dispersion. Specifically, we estimate a statistical model for normalized prices, in which both the store and the store-good component of prices are given by the sum of a fixed effect and an ARMA process. Our estimates imply that almost all of the variance of the store component is due to persistent differences in the average price of the store. Moreover, approximately half of the variance of the store-good component is due to persistent differences in the relative price of the good across stores, and half is due to transitory differences. Overall, a sizeable fraction of the variance of the price at which the same good is sold in the same period of time and in the same geographical area is due to persistent differences across stores in the price of that good relative to the price of the store.[2]

Price theory offers compelling explanations for the existence of differences in the average price of goods across different stores (i.e. the dispersion in the store component of prices), as well as for the existence of transitory fluctuations in the price of a particular good at a particular store (i.e. the dispersion in the transitory part of the store-good component of prices). Dispersion in the store component of prices is typically attributed to amenities. Indeed, if different stores offer different amenities to their customers and the production cost of these amenities is added to the price of goods, there will be variation in the average price of different stores. Dispersion in the transitory part of the store-good component of prices is typically attributed to intertemporal price discrimination (see, e.g., Conlisk, Gerstner and Sobel 1984, Sobel 1984 and Albrecht, Postel-Vinay and Vroman 2013). Indeed, stores can lower the price of a good for a short period of time in order to discriminate between low-valuation customers who are able to substitute their purchases intertemporally, and high-valuation customers who need to make their purchase at a particular time. However, price

[2]Sorensen (2000) was the first to document within store price dispersion in the retail market for medical drugs. In particular, he showed that the difference in the average price of different pharmacies accounted for only a fraction of the overall cross-sectional variation in the prices at which a particular drug was sold. However, since Sorensen did not follow the pharmacy over time, he could not distinguish between transitory and persistent differences in relative prices. Kaplan and Menzio (2014b) tried to distinguish between transitory and persistent differences in relative prices. However, they used the Kilts Nielsen Consumer Panel Dataset which contains much fewer price observations than the KNRS.

theory does not offer much guidance in the way of understanding the existence of permanent differences across stores in the price of a particular good relative to the average price of the store. Presumably, this gap in the literature is partly due to the fact that permanent dispersion in relative prices had not been documented before, and it is partly due to the fact that a theory of relative price dispersion requires developing an equilibrium model of a retail market where sellers trade multiple products.[3]

In the second part of the paper, we propose a theory of relative price dispersion. We consider an imperfectly competitive retail market in which sellers and buyers respectively supply and demand two goods. Sellers are ex-ante homogeneous, both with respect to their cost of producing the goods and in terms of the type of buyers they meet. Buyers are ex-ante heterogeneous. One type of buyers, which we call busy, have a relatively high valuation for goods and they need to make all of their purchases in the same store (even if they have access multiple sellers). The other type of buyers, which we call cool, have a relatively low valuation and they can purchase different items at different stores (as long as they have access to multiple sellers). The retail market is imperfectly competitive. As in Butters (1977) or Burdett and Judd (1983), a buyer does not have access to all the sellers in the market. Instead, a buyer has access to just one seller with some positive probability and to multiple sellers with complementary probability. As it is well-know from Butters (1977), Varian (1980) and Burdett and Judd (1983), this structure of contacts between buyers and sellers generates a distribution of prices that lie in between the competitive and the monopoly levels. Moreover, as the fraction of buyers in contact with multiple sellers goes to one, all prices converge to the competitive level. In contrast, as the fraction of buyers in contact with multiple sellers goes to zero, all prices converge to the monopoly level.

We show that, for some parameter values, the equilibrium must feature relative price dispersion. To this aim, we first show that some sellers find it optimal to post different prices for the two goods. Indeed, consider a seller that sets the same price for the two goods and suppose that this common price lies in between the valuation of the cool buyers and the valuation of the busy buyers. Given this common price, the seller trades with some of the

[3]The theoretical literature studying the pricing problem of multiproduct sellers in a market with search frictions is rather thin. McAfee (1995) studies a multiproduct version of Burdett and Judd (1983) with recall. Baughman and Burdett (2015) study a version of Burdett and Judd (1983) without recall. More recently, Zhou (2014) and Rhodes (2015) study the price setting problem of multi-product sellers. However, none of these papers considers the type of heterogeneity among buyers and the resulting price discrimination behavior of sellers that are at the core of our theory. There is also a literature that, in the context of the Hotelling (1929) model of imperfect competition studies the pricing problem of multi-product sellers (see, e.g., Lal and Matutes 1994, Ellison 2005). This literature has been mostly concerned with add-on pricing, i.e. optimal pricing when buyers know some but not all of the prices posted by a seller.

busy buyers it encounters, but it never trades with the cool buyers. Now, suppose that the seller lowers the price of the first good and increases the price of the second good, so as to keep the average price constant. Since busy buyers must purchase both goods in the same location, their utility from purchasing from the seller is unchanged. Hence, the seller makes the same number of trades with the busy buyers. However, as the price of the first good falls below the valuation of the cool buyers, the seller can also trade this one good to the cool buyers. Overall, the seller is strictly better off setting different prices for the two goods than setting a common price. Next, we show that the equilibrium is symmetric. That is, for every seller that posts a lower price for the first good than for the second good, there is another seller than posts a lower price for the second good than for the first good. As a result, relative price dispersion emerges in equilibrium.

The key to our theory of relative price dispersion is a form of discrimination between different types of buyers. The difference in valuation between the two types of buyers gives seller a motive to try to price discriminate. The difference in the ability to make purchases at multiple locations gives sellers a way to price discriminate. Price discrimination takes the form of an asymmetric pricing strategy which, in general equilibrium, leads to relative price dispersion.[4]

In the last part of the paper, we provide some evidence that is consistent with our theory of relative price dispersion. We use the Kilts-Nielsen Consumer Panel Dataset (KNCP), which tracks the shopping behavior of approximately 50,000 households over the period 2004-2009. Using the KNCP, we show that there is variation in the number of stores at which different households shop in a given month. The majority of households make all of their purchases in a single store, but a sizeable fraction purchase from multiple stores as well. Second, following the same methodology as in Aguiar and Hurst (2007), we show that there is a great deal of variation in the price index of different households. Finally, we show that households who shop with multiple sellers tend to pay lower prices, in the sense that their price index is lower than for households who do all of their shopping with the same seller. These observations show that, as assumed by our theory, some households do their

[4] After developing our theory, we discovered a predecessor in Lal and Matutes (1989). They consider a Hotelling (1929) model in which two sellers trade two goods to a population of buyers that are heterogeneous with respect to their location, their valuation and their commuting cost. They show that, for some parameter values, there is relative price dispersion and that relative price dispersion is used as a discrimination device. Even though the spirit of the two theories is similar, there are many modeling and substantive differences. For example, Lal and Matutes consider a Hotelling model of imperfect competition, while we study a model of imperfect competition a la Burdett and Judd (1983). We find our choice natural, as Burdett and Judd is the canonical model for studying equilibrium price dispersion. Moreover, while the Hotelling model does not always admit an equilibrium (which, indeed, is the case in Lal and Matutes), the Burdett and Judd model does not suffer from this type of problem.

shopping in one location and some in multiple locations and that, as predicted by our theory, the households who shop in multiple locations end up paying lower prices.

2 Relative Price Dispersion: Evidence

In this section, we document the extent and analyze the sources of the dispersion in the price at which the same good is sold in the same period of time and in the same geographical area. We use a rich dataset on prices that includes the price of multiple goods at each store, and the same store over time. With these data, we estimate a rich stochastic process for the average price level of a store, as well as for the price of a good at a store relative to the average price level of the store. We use the estimates of the stochastic process to decompose the variance in the price of the same good in the same period of time and in the same area. The main finding is that a significant fraction of the cross-sectional price variance is due to the fact that stores that are on average equally expensive set persistently different prices for the same good. We refer to this phenomenon as relative price dispersion.

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 defined to be one week and a good is defined by its UPC (barcode). We first decompose the log of each price p_{jst} into three additively separable components: a component that reflects the average price of the good in period t; μ_{jt}; a component that reflects the expensiveness of the store selling the good, y_{st}; and a component that reflects factors that are unique to the combination of store and good, z_{jst}.[5]. Formally, we decompose the log of p_{jst} as

$$\log p_{jst} = \mu_{jt} + y_{st} + z_{jst} \tag{1}$$

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-covariance functions of y_{st} and z_{jst}, which are illustrated in Figure 1. The auto-covariance functions of y_{st} and z_{jst} display a sharp drop at short lags, followed by a smoothly declining profile that remains strictly

[5]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.

positive even at very long lags. The initial drop in the auto-covariance suggests the presence of a transitory component in both y_{st} and z_{jst}. We model the transitory components as a MA(q) process, rather than as in IID process, to allow for the possibility that the transitory component may reflect temporary sales. 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. The smoothly declining portion of the auto-covariance function is consistent with the presence of an AR(1) component. Finally, the fact that the auto covariance function remains positive even after 100 weeks suggests the presence of a fixed effect.

Figure 1: Autocorrelation function of prices

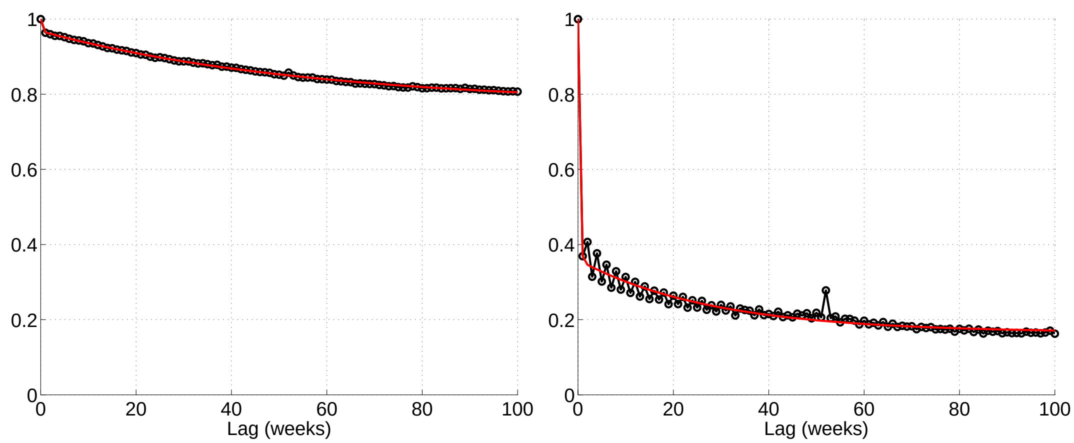

(a) Store component (b) Store-good component

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

Formally, the statistical model for y_{st} and z_{jst} is given by

$$
\begin{aligned}
y_{st} &= y^F_s + y^P_{st} + y^T_{st} & z_{jst} &= z^F_{js} + z^P_{jst} + z^T_{jst} \\
y^P_{st} &= \rho_y y^P_{s,t-1} + \eta^y_{s,t} & z^P_{jst} &= \rho_z z^P_{js,t-1} + \eta^z_{js,t} \\
y^T_{st} &= \varepsilon^y_{s,t} + \sum_{i=1}^{q} \theta_{y,i} \varepsilon^y_{s,t-i} & z^T_{jst} &= \varepsilon^z_{js,t} + \sum_{i=1}^{q} \theta_{z,i} \varepsilon^z_{js,t-i} \\
y^F_s &= \alpha^y_s & z^F_{js} &= \alpha^z_{js}
\end{aligned} \tag{2}
$$

where y^F_s and z^F_{js} denote the fixed-effects of the store and of the store-good components, y^P_{st} and z^P_{jst} denote the persistent parts of the store and of the store-good components, and y^T_{st} and z^T_{jst} denote the transitory parts of the store and of the store-good components. The parameters α^y_s and α^z_{js} are normal 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 correlation coefficients 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 normal 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. We experimented with alternative specifications of $\varepsilon^z_{js,t}$ which are meant to capture the possibility that the MA(q) part of the store-good component is due to sales and, hence, might be better described by a left-skewed distribution. However, our findings were substantively unchanged.

We estimate the parameters of the statistical model in (2) using data on quantity-weighted average prices, p_{jst}, for a large number of goods $j = 1 \ldots J$, at a large number of stores $s = 1 \ldots S$ in a single geographic market m at a weekly frequency $t = 1 \ldots T$. 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 infeasible. Instead we estimate the model using a multi-stage 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 and Blundell and Preston 1998). Notice that we assumed that the statistical model in (2) has the same parameters for every good. In a robustness check, we estimated (2) separately for different categories of goods and found rather similar results.

The estimation procedure involves four steps.

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

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

We then construct normalized prices as

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

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

$$\hat{y}_{st} = \frac{1}{n_{jst}}\sum_{j=1}^{n_{jst}} \tilde{p}_{jst} \tag{5}$$

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 then estimate the store-good component z_{jst} as

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

The process described above 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 Kilts-Nielsen Retail Scanner Dataset (KNRS). The KNRS 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 over 1.5 million UPCs at around 40,000 stores in over 2,500 counties across 205 Designated Market Areas (DMA). A DMA is a geographic area defined by Nielsen that is roughly the same size as a Metropolitan Statistical Area (MSA). 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 over 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 UPCs.

For concreteness, we start by focusing on a single DMA: Minneapolis-St Paul. We then show that our findings are robust to extending the analysis to cover a broad set of geographically dispersed markets. Our baseline set of UPCs for the Minneapolis-St Paul DMA is chosen as the 1000 UPCs with the largest quantities of sales in the state of Minnesota in the first quarter of 2010. These 1,000 products span 50 product groups. Table 7 in Appendix A shows how these 1000 sample UPCs are distributed across goods departments. Table 8 in Appendix A shows the number and percentage of UPCs in the 20 product groups with the highest representation among these 1000 UPCs. To give a sense of how frequently these

Table 1: Parameter estimates

Parameter	Baseline	State	County	$N_1 = 50$	$N_1 = 500$	$N_2 = 25$	$N_2 = 100$	UPC 1	UPC 2
	Store component								
ρ^y	0.983	0.983	0.991	0.987	0.990	0.983	0.984	0.869	0.989
θ	0.000	0.039	0.218	0.087	0.000	0.000	0.006	0.000	0.039
$Var(\alpha^y)$	0.003	0.003	0.002	0.005	0.002	0.003	0.003	0.005	0.003
$Var(\eta^y)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$Var(\epsilon^y)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	Store-good component								
ρ^z	0.965	0.965	0.964	0.970	0.965	0.965	0.964	0.967	0.967
ϕ	0.026	0.042	0.035	0.039	0.016	0.027	0.032	0.074	0.242
$Var(\alpha^z)$	0.003	0.004	0.003	0.005	0.003	0.003	0.003	0.005	0.003
$Var(\eta^z)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$Var(\epsilon^z)$	0.013	0.013	0.013	0.016	0.013	0.012	0.013	0.014	0.006

Notes: The baseline model is estimated on data for the Minneapolis-St Paul designated market area, and the next columns present results for data on the entire state of Minnesota, as well as Hennepin County alone. The next columns present results for different sample section criteria, including alternative sets of UPCs: UPC 1 refers to the 1463 UPCs used in the nationwide analysis, while UPC 2 refers to the alternative set of 100 UPCs described in the text.

products are purchased, in the Minneapolis market areas in 2010:Q1, the product with the largest quantity of units sold was in the Fresh Eggs product module, of which nearly 2.9 million units were sold. The least frequently sold of these 1000 products was in the Liquid Cocktail Mixes product module of which just under 50,000 units were sold.

Even after restricting attention to these 1000 products, the dataset is extremely large. Over the 7 year period from 2006 to 2012, we have over 40 million observations of prices p_{jst}. To ensure that our findings are not specific to this particular bundle of goods, we also estimate the model using two alternative sets of UPCs. First, we select the 1000 UPCs ranked 9001 to 10000 based on the aforementioned list from Minnesota. The idea to choose an alternative set of less frequently purchased products, for which there are still enough transactions for reliable estimation. Second, we select the set of UPCs that were either among the top 1000 most commonly purchased UPCs nationwide in 2010 based on quantity, or were among the top 1000 most commonly purchased UPCs nationwide in 2010 based on revenue. The resulting set of 1463 UPCs is the one we use when comparing results across different geographic areas.

We estimate the model separately for each geographic area. For a given set of UPCs and a given geographic area, we select stores, goods and weeks that satisfy two criteria:

1. For each store/week combination, we have quantity and price data for at least N_1 of the

UPCs in in the given set. In our baseline estimation we set $N_1 = 250$, and 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 we set $N_2 = 50$, and we report results for $N_2 \in \{25, 100\}$.

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

2.3 Estimation results and variance decomposition

We first present results for the Minneapolis-St Paul designated market area. We then consider the robustness of these results to a range of alternative specifications, including whether they vary significantly across markets in the US.

2.3.1 Minneapolis-St Paul

We start by presenting baseline estimates and robustness analyses for the Minneapolis-St Paul market area. Figure 1 displays the fit of the auto-correlation function for the store component (Panel A) and the store-good component of prices (Panel B) out to 100 lags. The parameter estimates that correspond to this model are reported in the first column of Table 1. For both components, the statistical model provides an excellent fit to the shape of the autocorrelation function. Several features of the autocorrelation functions are worthy of mention. First, the auto-correlation of the store component is above 0.8 even at long lags, foreshadowing our finding that almost all of the store component is persistent in nature. Second, the sharp drop in the auto-correlation of the store-good component after one lag suggests the presence of a large transitory component in prices. Third, the slow exponential decay and then flattening out of the store-good component suggest the presence also of a substantial persistent part of the store-good component. Fourth, the spike at 52 weeks reflects the fact that some products display annual regularities in their prices. Finally, the zig-zag pattern of the auto-correlation of the store-good component is due to regularities in the patterns of sales that cannot be captured by our statistical model.

Overall, the estimated model fits the data very well and, for this reason, we are comfortable using it to decompose the cross-sectional variance of the price at which the same good is sold in the same week and in the same market. The variance decomposition is reported in

Table 2: Dispersion in prices: persistent and transitory

	Variance		Percent		Std. Dev.
Store					
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					
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.

Table 2. The variance of the price of the same good in the same week and market is 0.023 or, equivalently, the standard deviation is 15.3%. The variance of the store component accounts for 15% of the overall variance of the price, and the variance of the store-good component accounts for the remaining 85%. That is, most of the variation in the price at which a good is sold is not due to the fact that the good is sold at stores that are, on average, more or less expensive. Most of the variation in the price at which a good is sold is due to the fact that the good is sold at different prices at stores that are, on average, equally expensive.

The variation in prices associated with the store and store-good components could be due to either the transitory or the permanent component. The statistical model (2) is designed to distinguish between these two sources of variation. Since the estimated persistence of the AR(1) component of prices is extremely close to unity for both the store and store-good components (the estimates of ρ_z and ρ_y for the baseline model are 0.965 and 0.983, respectively), we group the fixed effect and AR(1) components together and refer to these as the persistent part of the price, and we refer to the MA components as the transitory part of the price.

The decomposition in Table 2 reveals that nearly all of the price variance that is due to variation in the store component comes from persistent differences in the average price of different stores. In contrast, 65% of the price variance that is due to the variation in the store-good component comes from transitory differences across stores in the price of the good relative to the average price of the store. Yet, a sizeable fraction of the price variation that is due to the variation in the store-good component comes from persistent differences

Table 3: Robustness to geographic area

	Baseline/DMA Minneapolis-St Paul		State Minnesota		County Hennepin	
	Sd	Decomp/%	Sd	Decomp/%	Sd	Decomp/%
Store						
Transitory	0.011	3.2	0.011	3.8	0.015	6.2
Fixed plus Pers.	0.059	96.8	0.058	96.2	0.058	93.8
Total Store	0.060	15.5	0.059	14.4	0.060	15.9
Store-good						
Transitory	0.113	64.1	0.115	63.5	0.114	67.6
Fixed plus Pers.	0.084	35.9	0.087	36.5	0.079	32.4
Total Store-good	0.141	84.5	0.144	85.6	0.138	84.1
Total	0.153	100.0	0.155	100.0	0.151	100.0

Notes: This table presents a robustness exercise comparing our baseline results focusing on the Minneapolis-St Paul designated market area to results using data on the entire state of Minneapolis and data on Hennepin county alone.

across stores in the relative price of the good. This is what we call relative price dispersion. Relative price dispersion is a feature of the data that had not been well documented before and, at first blush, it seems hard to rationalize. Why would do stores that are on average equally expensive choose to systematically charge different prices for the very same good?

Finally, notice that while variance decompositions are a convenient tool for breaking down dispersion into orthogonal elements, the fact that variances are measured in squared prices makes the comparison of the various elements somewhat hard to interpret. For this reason, the final column of Table 2 reports the standard deviation of each of the orthogonal components of prices implied by the model. The overall standard deviation of prices is 15% and the standard deviation due to persistent differences in relative prices is 8%. These figures perhaps further emphasize the point that persistent differences in relative prices are an important feature of the retail market.

2.3.2 Robustness checks

The estimates in Table 2 highlight two important features of price dispersion, both of which turn out to be extremely robust. First, the vast majority of price dispersion is due to variation in the store-good component of prices, rather than to the store component of prices. Second, of the variation in the store-good component, at least one-third is due to highly persistent differences across stores in the price of the good relative to the price of the store. Tables 1, 3, 4, and 5 report the parameter estimates and the variance decomposition for various

Table 4: Robustness to sample criteria

	Baseline		$N_1 = 50$		$N_1 = 500$		$N_2 = 25$		$N_2 = 100$	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store										
Transitory	0.011	3.2	0.017	5.0	0.011	3.6	0.011	3.2	0.011	3.3
Fixed plus Pers.	0.059	96.8	0.075	95.0	0.055	96.4	0.058	96.8	0.063	96.7
Total Store	0.060	15.5	0.077	19.0	0.056	13.6	0.059	15.3	0.064	16.8
Store-good										
Transitory	0.113	64.1	0.126	63.3	0.113	65.4	0.111	63.8	0.114	65.0
Fixed plus Pers.	0.084	35.9	0.096	36.7	0.082	34.6	0.084	36.2	0.084	35.0
Total Store-good	0.141	84.5	0.158	81.0	0.140	86.4	0.139	84.7	0.141	83.2
Total	0.153	100.0	0.176	100.0	0.151	100.0	0.151	100.0	0.155	100.0
	Baseline		Weighted		UPC 2		UPC 2 Weight		UPC 1	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store										
Transitory	0.011	3.2	0.019	21.0	0.011	2.2	0.019	19.3	0.006	1.4
Fixed plus Pers.	0.059	96.8	0.037	79.0	0.072	97.8	0.040	80.7	0.054	98.6
Total Store	0.060	15.5	0.041	6.5	0.073	18.9	0.044	6.2	0.055	18.7
Store-good										
Transitory	0.113	64.1	0.124	62.6	0.119	61.0	0.132	58.5	0.082	51.1
Fixed plus Pers.	0.084	35.9	0.096	37.4	0.095	39.0	0.111	41.5	0.080	48.9
Total Store-good	0.141	84.5	0.156	93.6	0.152	81.1	0.173	93.8	0.114	81.3
Total	0.153	100.0	0.162	100.0	0.169	100.0	0.178	100.0	0.127	100.0

Notes: This table presents a robustness exercise comparing our baseline results to results obtained using alternative sample section criteria: alternative cutoffs for required numbers of observations, quantity weighting in constructing the store and store-good components, and alternative sets of UPCs: UPC 1 refers to the 1463 UPCs used in the nationwide analysis, while UPC 2 refers to the alternative set of 100 UPCs described in the text.

alternative cuts of the data.

First, we consider alternative levels of geographic aggregation for the definition of a market. In Table 3 we report the variance decomposition when we use a broader definition of market (i.e., Minnesota state) and a narrower definition of a market (i.e., Hennepin county) than in the baseline (i.e., Minneapolis-Saint Paul). The reader can clearly see that the variance decomposition is essentially unchanged across the three market definitions.

Second, we consider alternative selection criteria for the minimum number of goods sold for a store/week to be included in the sample (N_1), and the minimum number of stores for a good/week to be included in the sample (N_2). The variance decompositions are shown in Table 4 and they are very similar to the baseline decomposition in Table 2.

Third, we consider alternative samples of UPCs and the effects of using quantity-weighted, rather than raw averages in our construction of sample moments. The decompositions from these alternative samples are also shown in Table 4. Both moving to the broader sets of UPCs and using quantity-weighting leads to an increase in the fraction of the overall price

Table 5: Robustness to statistical model and estimation weights

	Baseline		Identity Weight		MA(5)	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store-good						
Transitory	0.113	64.1	0.113	64.4	0.114	65.6
Fixed plus Pers.	0.084	35.9	0.084	35.6	0.082	34.4
Total Store-good	0.141	100.0	0.141	100.0	0.141	100.0
	Skewed MA(1)		**Uniform Sales**		**Identity Weight**	
	Sd	Dec/%	Sd	Dec/%	Sd	Dec/%
Store-good						
Transitory	0.113	64.1	0.098	57.7	0.124	62.6
Fixed plus Pers.	0.084	35.9	0.084	42.3	0.096	37.4
Total Store-good	0.141	100.0	0.129	100.0	0.156	100.0

Notes: This table presents a robustness exercise comparing our baseline results to results obtained using alternative GMM-estimation weights (unit weighting), extending the MA process to more lags (5), allowing for skewness in the MA innovations, and modeling the transitory variation with an explicit model of sales described in the appendix.

variation that can is due to persistent differences in the relative price of the good at different stores.

Finally, we consider the effects of using a different weighting matrix in the GMM estimation and alternative ways of modeling the transitory part of the store-good component. We consider allowing for an MA(5) rather than an MA(1), allowing for skewness in the transitory innovations, and replacing the MA process with an explicit model of uniformly distributed sales (see Appendix A). Our main findings are robust to all of these alternative specifications.

2.3.3 Nationwide estimates

The analysis in the previous section focused on a single geographic region. It is natural to ask whether our findings apply to other geographical areas. In this section we show that the insights from Minneapolis-St Paul extend to the remainder of the United States. 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 in our data sets and repeated the estimation for each market, using the same set of 1463 UPCs for each market. As described above, this set of UPCs was chosen to reflect UPCs that are commonly purchased nationwide.

The top panel of Figure 2 shows histograms of the standard deviation of prices in each of the 25 DMAs, as well as the fraction of the variance due to the store versus store-good components, and the fraction of the variance of each component that is due to transitory versus persistent factors. The analogous statistics for the 25 counties are displayed in the bottom panel of Figure 2.

Designated Market Areas

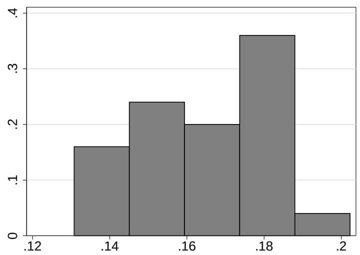

(c) St. dev. of UPC prices

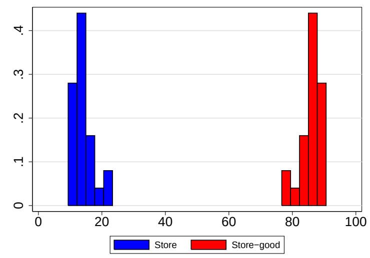

(d) Store vs. store-good component

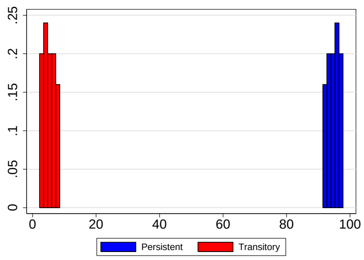

(e) Store: permanent vs. transitory component

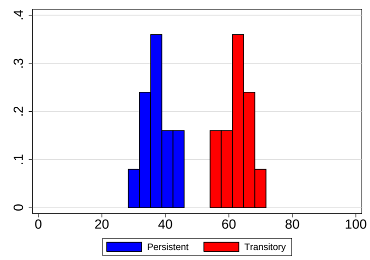

(f) Store-good: permanent vs. transitory component

Counties

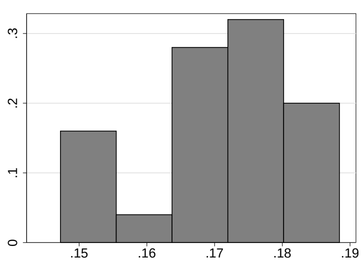

(g) St. Dev. of UPC prices

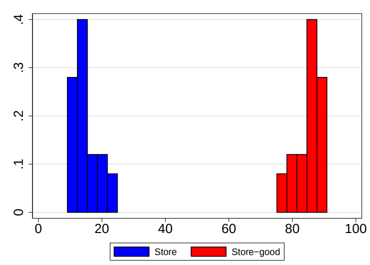

(h) Store vs. Store-good component

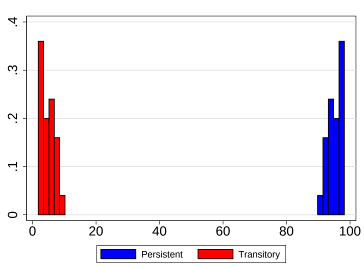

(i) Store: Permanent vs. transitory component

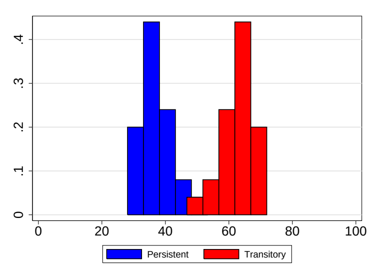

(j) Store-good: Permanent vs. transitory component

Notes: These figures present a robustness exercise comparing our baseline results to results obtained for alternative markets. The top panel presents histograms showing how the results vary across designated market areas in the US. The bottom panel presents histograms showing how the results vary across counties in the US.

Figure 2: Price dispersion and variance decompositions across geographic areas

These figures clearly show 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 prices occurs in 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.

3 Relative Price Dispersion: Theory

In the previous section, we documented the existence of persistent dispersion in relative prices across retailers operating in the same geographical area. In this section, we advance a theory of relative price dispersion in the context of the canonical model of imperfect competition of Burdett and Judd (1983). According to our theory, relative price dispersion does not emerge because of differences in the relative wholesale cost, or differences in the relative elasticity of demand, of different goods across retailers. Instead, we propose the view that relative price dispersion emerges because sellers want to price discriminate between high-valuation buyers who need to make all of their purchases in the same store, and low-valuation buyers who are willing to purchase different goods in different stores. In Section 4, we will provide some empirical evidence to support our theory.

3.1 Environment

We consider an imperfectly competitive retail market in which a population of homogeneous sellers sells two goods (i.e., good 1 and good 2) to a population of heterogeneous buyers.

On one side of the market, there is a measure 1 of sellers. Every seller is able to produce each one of the two goods at a marginal cost c, which we normalize to zero. Every seller chooses a price for good 1, p_1, and a price for good 2, p_2, taking as given the joint distribution of sellers over price pairs, $H(p_1, p_2)$, and the associated marginal distribution of sellers over the price of good 1, $F_1(p_1)$, and over the price of good 2, $F_2(p_2)$. Every seller chooses the prices (p_1, p_2) so as to maximize their profit.

On the other side of the market, there is a measure 1 of buyers. A fraction μ_b of the buyers are of type b and a fraction $\mu_c = 1 - \mu_b$ are of type c, where b is mnemonic for *busy*, c is mnemonic for *cool* and $\mu_b \in (0, 1)$. Every busy buyer demands one unit of each good, for which he has valuation $u_b > 0$. Every cool buyer demands one unit of each good, for which he has valuation $u_c > 0$. Hence, if a buyer of type $i = \{b, c\}$ purchases both goods at the prices (p_1, p_2), he attains a payoff of $2u_i - p_1 - p_2$. If a buyer of type i purchases only good $k = \{1, 2\}$ at the price p_k, he attains a payoff of $u_i - p_k$. Finally, if a buyer does not purchase either good, he attains a payoff of zero.

Competition in the retail market is imperfect. We assume that a buyer cannot purchase from any seller in the market, as they are in contact with only a subset of sellers. In particular, a fraction α_b of busy buyers is in contact with only one seller, while a fraction $1 - \alpha_b$ is in contact with multiple sellers, where $\alpha_b \in (0, 1)$. Similarly, a fraction α_c of cool buyers is in contact with only one seller, while a fraction $1 - \alpha_c$ is in contact with multiple sellers, where $\alpha_c \in (0, 1)$. We like to interpret these contacts as the set of sellers that are physically close to the buyer when he has to make a purchase. We refer to buyers who are in contact with only one seller as captive, and to buyers who are in contact with multiple sellers as non-captive. For the sake of exposition, we assume that non-captive buyers are in contact with two sellers.

As well explained in Butters (1977), Varian (1980) and Burdett and Judd (1983), the fraction of buyers who are in contact with multiple sellers determines the competitiveness of the retail market. Indeed, if $\alpha_b = \alpha_c = 0$, every buyer is in contact with multiple sellers and this is enough to guarantee than the retail market is perfectly competitive. If $\alpha_b = \alpha_c = 1$, every buyer is in contact with only one seller and the retail market is monopolistic. If α_b and/or α_c are between 0 and 1, the retail market is between competitive and monopolistic.

Busy buyers and cool buyers differ along two dimensions. First, we assume that busy buyers have a higher valuation for goods than cool buyers, i.e. $u_b > u_c$. Second, we assume that busy buyers must make all their purchases from the same seller, while cool buyers can make purchases from different sellers (among those with whom they are in contact). That is, a busy buyer may be in contact with one or two sellers and may purchase one or two goods, but he must make all of his purchases from one retailer. In contrast, if a cool buyer is in contact with multiple sellers, he can purchase good 1 from one retailer and good 2 from a different retailer. Both differences between busy and cool buyers can be seen as consequences of differences in wages. If busy buyers earn higher wages in the labor market, they will tend to have a higher valuation for goods than cool buyers. Similarly, if busy buyers earn higher wages in the labor market, they will tend to have a higher value of time than cool buyers. Since going from store to store to purchase different items is time consuming, busy buyers will prefer to purchase everything in the same place while cool buyers will be willing to purchase different items in different places. The differences between busy and cool buyers give sellers an incentive and an opportunity to price discriminate and, as we shall see, price discrimination will take the form of relative price dispersion.

It may be natural to think that busy buyers, as individuals with a higher opportunity cost of time, are also less likely to be in contact with multiple sellers than cool buyers.[6] However, this additional difference between the two types of buyers is not necessary for our theory of relative price dispersion. Therefore, in order to keep the exposition as simple as possible, we will assume that $\alpha_b = \alpha_c = \alpha$. Our results easily generalize to the case in which $\alpha_b > \alpha_c$.

A few comments about the environment are in order. First, we consider a retail market where two goods are traded. This is the simplest version of a retail market for which we can meaningfully talk about relative price dispersion, i.e. across-retailer variation in the price at which one good is sold at a store relative to the average price of goods at that store. Second, we assume that sellers have the same cost of production and face the same population of buyers. We make this assumption because we want to develop a theory of relative price dispersion that does not emerge simply from sellers facing different relative costs for the goods, or different relative elasticities of demand for the goods.[7] Finally, we consider a retail market that is static. Even though our empirical analysis of prices was dynamic, here we are interested in explaining the persistent component of prices and, hence, a dynamic model would simply introduce unnecessary complications.[8]

3.2 General properties of equilibrium

In this section, we establish some general properties of equilibrium. First, we identify the region in the $\{p_1, p_2\}$ space where sellers who post prices summing up to more than $u_b + u_c$ may find it optimal to locate themselves. Second, we identify the region in the $\{p_1, p_2\}$ space where sellers who post prices summing up to more than $2u_c$ but to less than $u_b + u_c$

[6]Indeed, this is the approach taken by Kaplan and Menzio (2014a), where buyers who are unemployed are assumed to spend more time searching the retail market and, hence, to be more likely to be in contact with multiple sellers than buyers who are employed.

[7]These explanations of relative price dispersion would be basically unfalsifiable, as data on wholesale costs and demand curves faced by different retailers is generally unavailable. Moreover, our prior is that retailers operating in the same market and in the same period of time are likely to face rather similar costs and demand. It is also useful to draw a parallel with the theoretical literature on temporary sales. Clearly, one could explain temporary sales with temporary declines in wholesale costs or with temporary increases in the elasticity of demand faced by retailers. However, the literature has tried to explain temporary sales in models without shocks.

[8]If we were to add dynamics to the model, we could also attempt to explain the transitory component of price dispersion, i.e. the variation in the price of a same good that is caused by the high-frequency variation in the price at which the same good is sold at the same retailer. Menzio and Trachter (2015) develop a dynamic model that is similar to the one presented here in which sellers post different prices on different days, as a way to discriminate between high-valuation buyers who can shop only on particular days, and low-valuation buyers who can shop every day.

find it optimal to locate themselves. Finally, we establish some properties of the marginal distribution of sellers over basket prices and some properties of the seller's equilibrium profit. To carry out the analysis we find it useful to denote as $G(q)$ the fraction of sellers whose prices (p_1, p_2) sum up to less than q, and with $\nu(q)$ the measure of sellers whose prices sum up to equal to q. Similarly, we denote as $F_i(p)$ the fraction of sellers whose price for good $i = \{1, 2\}$ is less than p, and with $\lambda_i(p)$ the measure of sellers whose price for good i is equal to p.

Figure 3: Pricing decision of sellers

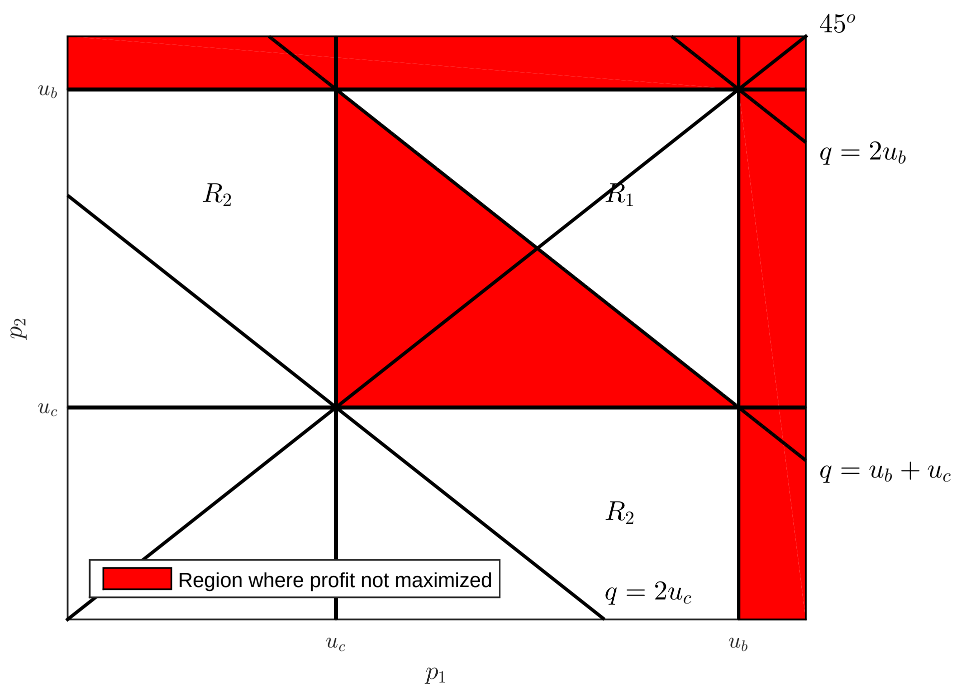

Notes: This figure considers the pricing decision of sellers discussed in the text, illustrating which regions of the (p_1, p_2) space will not be profit-maximizing. Conditional on a basket price in the interval $u_b + u_f \leq q \leq 2u_b$, sellers will not price outside R_1. Conditional on a basket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

Our first lemma shows that, if there is a seller who posts prices (p_1, p_2) such that the price $q = p_1 + p_2$ of the basket of goods is strictly greater than $u_b + u_c$, then this seller must be located in the triangular region R_1 in Figure 3. That is, this seller must set prices (p_1, p_2) that are each in between the valuation of cool buyers u_c and the valuation of busy buyers u_b. This property of equilibrium follows from the fact that a seller never finds it optimal to post a price greater than the valuation of busy buyers u_b.

Lemma 1: (i) A seller never finds it optimal to post prices (p_1, p_2), where either p_1, p_2 or both

p_1 and p_2 are strictly greater than u_b. (ii) If a seller post prices (p_1, p_2) with $p_1+p_2 > u_b+u_c$, then (p_1, p_2) belongs to the set

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

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

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

$$P(p_1, p_2) = \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2\right]\right\} p_2. \tag{8}$$

The expression in (8) is easy to understand. The seller is contacted by $\mu_b \alpha$ captive buyers of type b. A captive buyer of type b purchases good 2 from the seller with probability 1. The seller is also contacted by $2\mu_b(1-\alpha)$ non-captive buyers of type b. A non-captive buyer of type b purchases good 2 from the seller with probability $1 - \hat{G}(u_b + p_2) - \hat{\nu}(u_b + p_2)/2$, where $1 - \hat{G}(u_b + p_2)$ denotes the fraction of sellers that charge prices (p_1', p_2') such that $\min\{p_1', u_b\} + \min\{p_2', u_b\} > u_b + p_2$, and $\hat{\nu}(u_b + p_2)$ is the measure of sellers that charge prices (p_1', p_2') such that $\min\{p_1', u_b\} + \min\{p_2', u_b\} = u_b + p_2$. The seller is also contacted by $\mu_c \alpha$ captive buyers of type c and by $2\mu_c(1-\alpha)$ non-captive buyers of type c, but it does not trade with any of them as both of its prices are greater than u_c. If the seller instead posts the prices (p_1', p_2') such that $p_1' = u_b$ and $p_2' = p_2$, it attains a profit of

$$\begin{aligned} P(p_1', p_2') &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - \hat{G}(p_1' + p_2') + \hat{\nu}(p_1' + p_2')/2\right]\right\}(p_1' + p_2') \\ &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2\right]\right\}(u_b + p_2) \\ &> P(p_1, p_2). \end{aligned} \tag{9}$$

The strict inequality in (9) follows from the fact that, by lowering the price of good 1 to u_b, the seller trades with the same probability with buyers of type b, but it sells to them both good 1 and good 2, rather than only good 2. Hence, a seller never finds it optimal to post prices (p_1, p_2) with $p_1 > u_b$ and $p_2 \in (u_c, u_b]$. For the very same reason, a seller never finds it optimal to post prices (p_1, p_2) with $p_1 \in (u_c, u_b]$ and $p_2 > u_b$.

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

$$\begin{aligned} P(p_1,p_2) \quad &= \mu_b \left\{ \alpha + 2(1-\alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} p_2 \\ &+ \mu_c \left\{ \alpha + 2(1-\alpha) \left[1 - F_2(p_2) - \lambda_2(p_2)/2 \right] \right\} p_2. \end{aligned} \tag{10}$$

The expression in (10) is the same as in (8) with the addition of the term in the second line. This term represents the profit that the seller makes from trading good 2 to buyers of type c. If the seller instead posts the prices (p_1', p_2') with $p_1' = u_b$ and $p_2' = p_2$, it attains a profit of

$$\begin{aligned} P(p_1',p_2') \quad &= \mu_b \left\{ \alpha + 2(1-\alpha) \left[1 - \hat{G}(p_1' + p_2') + \hat{\nu}(p_1' + p_2')/2 \right] \right\} (p_1' + p_2') \\ &+ \mu_c \left\{ \alpha + 2(1-\alpha) \left[1 - F_2(p_2') - \lambda_2(p_2')/2 \right] \right\} p_2' \\ &= \mu_b \left\{ \alpha + 2(1-\alpha) \left[1 - \hat{G}(u_b + p_2) + \hat{\nu}(u_b + p_2)/2 \right] \right\} (u_b + p_2) \\ &+ \mu_c \left\{ \alpha + 2(1-\alpha) \left[1 - F_2(p_2) - \lambda_2(p_2)/2 \right] \right\} p_2 \\ &> P(p_1,p_2) \end{aligned} \tag{11}$$

The strict inequality in (11) follows from the fact that, by lowering the price of good 1 to u_b, the seller trades with the same number of buyers of type b, but it sells to them both goods 1 and good 2, rather than only good 2. Hence, a seller never finds it optimal to post prices (p_1, p_2) such that $p_1 > u_b$ and $p_2 \in [0, u_c]$. For the very same reason, a seller never finds it optimal to post prices (p_1, p_2) such that $p_1 \in [0, u_c]$ and $p_2 > u_b$.

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

The second lemma shows that, in region R_1, the seller is indifferent between any pair of prices (p_1, p_2) that is associated with the same basket price $q = p_1 + p_2$. This property of equilibrium follows directly from the fact that, in region R_1, the seller only trades with buyers of type b who, since they have to make all of their purchases in the same place, only care about the price of the basket of goods and not about the price of individual items.

Lemma 2: Fix any $q \in (u_b + u_c, 2u_b]$. The seller attains the same profit for all pairs of prices $(p_1, p_2) \in R_1$ such that $p_1 + p_2 = q$.

Proof: Fix any $q \in (u_b + u_c, 2u_b]$. Consider a seller posting the prices $(p_1, p_2) \in R_1$ with $p_1 + p_2 = q$. The seller attains a profit of

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

The seller is contacted by $\mu_b \alpha$ captive buyers of type b. A captive buyer of type b purchases both goods from the seller, as $p_1 \leq u_b$ and $p_2 \leq u_b$. The seller is also contacted by $2\mu_b(1-\alpha)$

non-captive buyers of type b. A non-captive buyer of type b purchases both goods from the seller with probability $1 - G(p_1 + p_2) - \nu(p_1 + p_2)/2$, where $1 - G(p_1 + p_2)$ is the fraction of sellers with a basket price greater than $p_1 + p_2$, and $\nu(p_1 + p_2)$ is the measure of sellers with a basket price equal to $p_1 + p_2$. Notice that the buyer's purchasing decision is based on comparing basket prices, rather than $\min\{u_b, p_1\} + \min\{u_b, p_2\}$, as we know from Lemma 1 that every seller posts prices below u_b. The lemma then follows immediately from the observation that $P(p_1, p_2)$ does not depend on p_1 and p_2 separately, but only on their sum. ■

The next lemma shows that, if there is a seller who posts prices (p_1, p_2) such that the price $q = p_1 + p_2$ of the basket is strictly greater than $2u_c$ and smaller than $u_b + u_c$, then the seller must be located in one of the two triangular regions R_2 in Figure 3. That is, the seller will never choose two prices that are both between u_c and u_b. Instead, the seller will always choose the price of one good to be lower than the valuation of a type-c buyer, and the price of the other good to be greater than the valuation of a type-c buyer. The optimality of this asymmetric pricing strategy is at the core of our theory of relative price dispersion.

Lemma 3: Consider an equilibrium in which there is a positive measure of sellers posting prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$. The prices (p_1, p_2) posted by each one of these sellers belong to the set

$$\begin{aligned} R_2 &= \{(p_1, p_2) : p_1 \in [0, u_c], p_2 \in (u_c, u_b], p_1 + p_2 \in (2u_c, u_b + u_c]\} \\ &\cup \{(p_1, p_2) : p_2 \in [0, u_c], p_1 \in (u_c, u_b], p_1 + p_2 \in (2u_c, u_b + u_c]\}. \end{aligned} \tag{12}$$

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

$$P(p_1, p_2) = \mu_b \{\alpha + 2(1 - \alpha) [1 - G(p_1 + p_2) + \nu(p_1 + p_2)/2]\} (p_1 + p_2).$$

The above expression is easy to understand. The seller is contacted by $\mu_b \alpha$ captive buyers of type b. A captive buyer of type b purchases both goods from the seller with probability one. The seller is contacted by $2\mu_b(1 - \alpha)$ non-captive buyers of type b. A non-captive buyer of type b purchases both goods from the seller with probability $1 - G(p_1 + p_2) + \nu(p_1 + p_2)/2$. The seller is also contacted by $\mu_c \alpha$ captive buyers of type c and by $2\mu_c(1 - \alpha)$ non-captive buyers of type c. A buyer of type c does not purchase any good from the seller as both p_1 and p_2 are higher than his valuation u_c.

If the seller deviates and posts the prices (p'_1, p'_2) with $p'_1 = u_c$, $p'_2 = p_1 + p_2 - p'_1$, the seller attains a profit of

$$\begin{aligned} P(p_1', p_2') &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - G(p_1'+p_2') + \nu(p_1'+p_2')/2\right]\right\}(p_1'+p_2') \\ &+\mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F_1(p_1') + \lambda_1(p_1')/2\right]\right\} p_1' \\ &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - G(p_1+p_2) + \nu(p_1+p_2)/2\right]\right\}(p_1+p_2) \\ &+\mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F_1(u_c) + \lambda_1(u_c)/2\right]\right\} u_c \\ &> P(p_1, p_2). \end{aligned}$$

The price p_1' is set equal to the valuation of buyers of type c and, hence, it is smaller than the valuation of buyers of type b, i.e.$p_1' = u_c < u_b$. The price p_2' is such that the price of the seller's basket is unchanged, i.e. $p_2' = p_1 + p_2 - p_1'$, and hence it is smaller than the valuation of buyers of type b, i.e. $p_2' \leq u_b$. Since the price of the basket is unchanged and both goods are priced below u_b, the seller sells both goods to the buyers of type b with the same probability as before and at the same profit margin. However, since $p_1 \leq u_c$, the seller also sells good 1 to some buyers of type c. In particular, it sells the good to a captive buyer of type c with probability 1, and to a non-captive buyer of type c with probability $1 - F_1(u_c) - \lambda_1(u_c)/2$. Overall, by changing its prices from (p_1, p_2) to (p_1', p_2'), the seller strictly increases its profit. This implies that there is no equilibrium in which a seller who posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$ sets $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$.

The above observation, combined with part (i) of Lemma 1, implies that, if a seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, it must be setting one of the two prices below the valuation of the buyers of type c, and the other price in between the valuation of the buyers of type c and the buyers of type b. ■

Lemma 3 is the key to our theory of relative price dispersion. The lemma states that, if competition induces some sellers to set the price q of the basket of goods between $2u_c$ and $u_b + u_c$, then these sellers will never set the same price for all goods. Indeed, if one of these sellers sets the same price $q/2$ for both goods, it only trades with buyers of type b. However, if the seller lowers one price below u_c and raises the other so as to keep the price of the basket constant, it makes some trades with some buyers of type c, without losing any revenues on the buyers of type b.

In the next lemma, we establish two additional results. In the first part of the lemma, we show that, in any equilibrium, the marginal distribution of sellers over basket prices, G, does not have mass points. This property of equilibrium obtains for the same reason as in Butters (1977), Varian (1980) or Burdett and Judd (1983). Specifically, if there was a mass point at q_0, a seller posting (p_1, p_2) such that $p_1 + p_2 = q_0$ could lower one of the two prices by an arbitrarily small amount and, instead of just selling to half of the contacted buyers of type b who are in touch with another retailer charging q_0, it could sell to all of them. Moreover, if there was a mass point at $q_0 = 0$, a seller posting $p_1 = p_2 = 0$ could raise its prices to u_b and, instead of attaining a profit of zero, it could sell at a positive price to all the contacted buyers of type b who are not in contact with any other retailer. In the second part of the lemma, we show that, in any equilibrium, the marginal distribution of sellers over the price of good $i = \{1, 2\}$, F_i, does not have any mass points for $p \in (0, u_c]$. The logic behind this result is also similar to Butters (1977), Varian (1980) and Burdett and Judd (1983). However, in the case of the distribution of prices for an individual good, we cannot rule out the possibility of a mass point for $p > u_c$ or for $p = 0$. There might be a mass point for $p > u_c$ because, when the price of a particular good is higher than the valuation of the buyer of type c, the seller never trades the good in isolation and its price does not play an allocative role. There might be a mass point at $p = 0$ because the fact that a seller trades one good at a price of zero does not imply that the seller's profit is zero.

Lemma 4: (i) In any equilibrium, the marginal distribution of sellers over basket prices, G, does not have any mass points. (ii) In any equilibrium, the marginal distribution of sellers over the price of good $i = \{1, 2\}$, F_i, does not have any mass points for all $p \in (0, u_c]$.

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

$$\begin{aligned} P(p_1, p_2) \; &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - G(\hat{q}) + \nu(\hat{q})/2\right]\right\} \hat{q} \\ &+ \textstyle\sum_{i=1}^{2} \mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F_i(p_i) + \lambda_i(p_i)/2\right]\right\} 1[p_i \leq u_c] p_i, \end{aligned}$$

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

$$\begin{aligned} P(p_1', p_2') \; &= \mu_b \left\{\alpha + 2(1-\alpha)\left[1 - G(\hat{q} - \epsilon)\right]\right\} (\hat{q} - \epsilon) \\ &+ \textstyle\sum_{i=1}^{2} \mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F_i(p_i - \epsilon_i) + \lambda_i(p_i - \epsilon_i)/2\right]\right\} 1[p_i \leq u_c](p_i - \epsilon_i) \\ &> P(p_1, p_2), \end{aligned}$$

where the inequality follows from the fact that $1 - G(\hat{q} - \epsilon) > 1 - G(\hat{q}) + \nu(\hat{q})$ and ϵ, ϵ_1 and ϵ_2 are all arbitrarily small. Since $P(p_1', p_2') > P(p_1, p_2)$, there cannot be a mass point at $\hat{q}$.

(ii) On the way to a contradiction, suppose there is an equilibrium where F_1 has a mass point at $\hat{p}_1 \in (0, u_c]$. If in equilibrium a seller posts the price $p_1 = \hat{p}_1$ for the first good, it must post a price $p_2 \in [0, u_b]$ for the second good. This seller attains a profit of

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

If the seller deviates and posts prices (p_1', p_2') with $p_1' = p_1 - \epsilon$, $p_2' = p_2$, for $\epsilon > 0$ arbitrarily small, it attains a profit of

$$\begin{aligned} P(p_1', p_2') \quad &= \mu_b \{\alpha + 2(1-\alpha)[1 - G(p_1 + p_2 - \epsilon)]\} (p_1 + p_2 - \epsilon) \\ &+ \mu_c \{\alpha + 2(1-\alpha)[1 - F_1(p_1 - \epsilon)]\} (p_1 - \epsilon) \\ &+ \mu_c \{\alpha + 2(1-\alpha)[1 - F_2(p_2) + \lambda_2(p_2)/2]\} 1[p_2 \leq u_c] p_i \\ &\geq P(p_1, p_2), \end{aligned}$$

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

The first part of Lemma 4 has two important implications. First, since G has no mass points, any equilibrium must feature dispersion in the price at which different retailers sell the same basket of goods. In the language of Section 2, any equilibrium features dispersion in average store prices. Second, since F has no mass points, any equilibrium must feature dispersion in the price at which different retailers sell a particular good. In the language of Section 2, any equilibrium features price dispersion. The second part of Lemma 4 will be useful to establish conditions under which an equilibrium features dispersion in the price of good 1 relative to the price of good 2 or, in the language of Section 2, relative price dispersion.

3.3 Bundled Equilibrium

In this section, we characterize the properties and derive the conditions for the existence of an equilibrium in which every seller in the market chooses to set a basket price greater than $u_b + u_c$, i.e. an equilibrium in which $G(u_b + u_c) = 0$. We denote this type of equilibrium as a Bundled Equilibrium for reasons that will be obvious shortly.

We already know several properties of a Bundled Equilibrium. First, Lemma 1 implies that, in a Bundled Equilibrium, every seller posts prices $(p_1, p_2) \in R_1$, i.e. prices such that $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. Second, Lemma 2 implies that, in a Bundled Equilibrium, every seller is indifferent between posting any pair of prices $(p_1, p_2) \in R_1$ that is associated with the same basket price q. Third, Lemma 4 implies that, in a Bundled Equilibrium, the marginal distribution of basket prices has no mass points. In light of these observations, the equilibrium profit for a seller posting prices $(p_1, p_2) \in R_1$, with $p_1 + p_2 = q$, can be written as

$$P_1(q) = \mu_b \{\alpha + 2(1-\alpha)[1 - G(q)]\} q. \tag{13}$$

Given the seller's profit function (13), we can characterize the properties of the equilibrium distribution of basket prices, G. The next lemma identifies the equilibrium value of the highest price on the support of G.

Lemma 5: In a Bundled Equilibrium, the highest basket price, $\overline{q}$, on the support of G is $2u_b$.

Proof: On the way to a contradiction, suppose there is an equilibrium with $\overline{q} < 2u_b$. A seller posting the basket price $\overline{q}$ attains a profit of $P_1(\overline{q}) = \mu_b \alpha \overline{q}$, where the expression for $P_1(\overline{q})$ follows from (13) and $G(\overline{q}) = 1$. If the seller deviates and posts the basket price $2u_b$, it attains a profit of $P_1(2u_b) = \mu_b \alpha 2u_b$, where the expression for $P_1(2u_b)$ follows from (13) and $G(2u_b) = 1$. Since $P_1(2u_b) > P_1(\overline{q})$, the seller does not find it optimal to post $\overline{q}$. Hence, there cannot be an equilibrium in which $\overline{q} < 2u_b$. Since Lemma 1 implies that a seller does not find it optimal to set a basket price $q > 2u_b$, there cannot be an equilibrium in which $\overline{q} > 2u_b$. Hence, in any equilibrium, $\overline{q} = 2u_b$. ■

The above lemma states that $\overline{q} = 2u_b$. This implies that, in a Bundled Equilibrium, the equilibrium profit of the seller, which we shall denote as P^*, is equal to the profit that is attained by a seller who sets the price $2u_b$ for the basket of goods and only trades with captive buyers of type b. That is, in an Bundled Equilibrium, P^* is given by

$$P^* = \mu \alpha 2u_b. \tag{14}$$

The next lemma shows that the support of the marginal distribution of basket prices has no gaps.

Lemma 6: In a Bundled Equilibrium, the support of G in some interval $[\underline{q}, \overline{q}]$.

Proof: On the way to a contradiction, suppose that there is an equilibrium in which the support of G is not an interval, i.e. there exists q_0 and q_1, with $0 \leq q_0 < q_1 \leq \overline{q}$ such that $G'(q_0) > 0$ and $G(q_0) = G(q_1)$. A seller posting a basket price of q_0, attains a profit of

$$P_1(q_0) = \mu_b \{\alpha + 2(1-\alpha)[1 - G(q_0)]\} q_0.$$

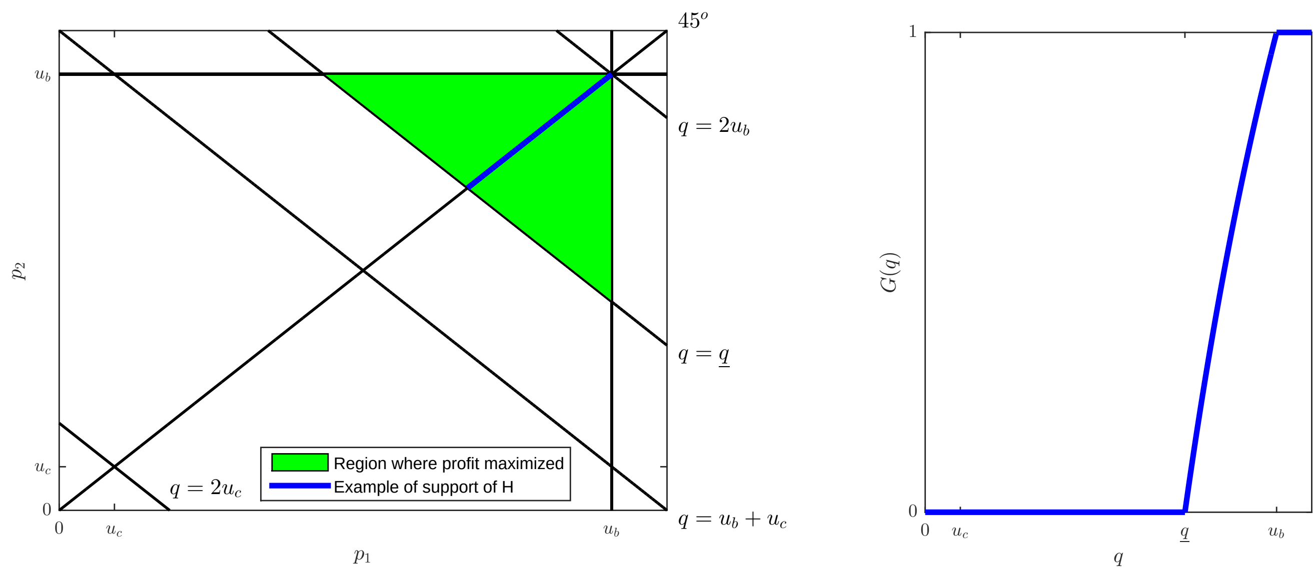

Notes: This figure shows the possible range of the support of the joint distribution $H(p_1, p_2)$, and the shape of the cumulative distributions $G(q)$, in the Bundled Equilibrium.

Figure 4: Bundled Equilibrium support of $H(p_1, p_2)$ and shape of $G(q)$

If the seller deviates and posts the basket price q_1, it attains a profit of

$$\begin{aligned} P_1(q_1) &= \mu_b \{\alpha + 2(1-\alpha)[1 - G(q_1)]\} q_1 \\ &= \mu_b \{\alpha + 2(1-\alpha)[1 - G(q_0)]\} q_1 > P_1(q_0). \end{aligned}$$

Since a seller never finds it optimal to set a basket price equal to q_0, q_0 cannot be on the support of G. Hence, there is no equilibrium where $G'(q_0) > 0$ and $G(q_0) = G(q_1)$. ∎

The previous lemma states that the support of the marginal distribution of basket prices is some interval $[\underline{q}, \overline{q}]$ with $\overline{q} = 2u_b$. Clearly, for every q on the support of the distribution, the seller must attain the equilibrium profit P^*, i.e.

$$\mu_b \{\alpha + 2(1-\alpha)[1 - G(q)]\} q = P^*, \forall q \in [\underline{q}, \overline{q}]. \tag{15}$$

Using (14) to substitute out P^* in (15), we obtain a functional equation that can be solved for the marginal distribution of basket prices. In particular, we find that G is given by

$$G(q) = 1 - \frac{\alpha}{2(1-\alpha)} \frac{2u_b - q}{q}. \tag{16}$$

Using (16) to solve the equation $G(q) = 0$ with respect to q, we find $\underline{q}$, the lower bound on the support of the marginal distribution of basket prices. In particular, we find that $\underline{q}$ is given by

$$\underline{q} = \frac{\alpha}{2-\alpha} 2u_b. \tag{17}$$

Figure 5 illustrates the key properties of a Bundled Equilibrium. In this type of equilibrium, the marginal distribution of basket prices is given by (16), which is exactly the same equilibrium price distribution as in the retail market with a single good that is studied by Burdett and Judd (1983). This is not surprising. In a Bundled Equilibrium, sellers only trade with buyers of type b, who, as they need to make all of their purchases from the same retailer, care exclusively about the price of the basket of goods and not about the price of each individual item. Therefore, in a Bundled Equilibrium, our model with two goods boils down to a version of Burdett and Judd in which the single item being traded is the basket of goods.

In a Bundled Equilibrium, the distribution G is pinned down and it is such that that there is dispersion in the price posted by different sellers for the same basket of goods, even though all sellers are identical with respect to the level of amenities offered to their customers (i.e., none) and all sellers are identical with respect to their marginal cost of production (i.e., zero). In the language of Section 2, the equilibrium features dispersion in store prices. In a Bundled Equilibrium, the joint distribution of sellers H over the price of the two goods is not uniquely pinned down. Indeed, any joint price distribution H that has support inside the region R_1 and that generates the marginal basket price distribution G in (16) is consistent with equilibrium. For example, there is an equilibrium where the joint price distribution is such that every seller with a basket price of q posts the same price, $q/2$, for both good 1 and good 2. In this equilibrium, there is no dispersion in the price of good 1 relative to the price of good 2 (or, equivalently, relative to the store price) across different sellers. However, there is also an equilibrium where a seller with a basket price of q posts a randomly chosen combination of prices $(p_1, p_2) \in R_1$ with $p_1 + p_2 = q$. In this equilibrium, there is dispersion in the price of good 1 relative to the price of good 2 (or, equivalently, relative to the store price) across different sellers. Therefore, there always exist a Bundled Equilibrium without relative price dispersion and a, fundamentally equivalent, equilibrium with relative price dispersion. Yet, in a Bundled Equilibrium, relative price dispersion is only a matter of indifference, as both buyers and sellers care exclusively about the price of the basket and not about the price of individual goods. In this sense, a Bundled Equilibrium does not provide a particularly interesting theory of relative price dispersion.

We conclude this subsection by identifying conditions for the existence of a Bundled Equilibrium. In particular, we can show that this type of equilibrium exists if and only if

$$\frac{\mu_c}{\mu_b} \leq \frac{3\alpha - 2}{(2-\alpha)u_c/u_b} - 1. \tag{18}$$

The above condition implies that a Bundled Equilibrium exists if and only if: (i) the product market is not too competitive, in the sense that α is greater than 2/3; and (ii) either the measure or the valuation of buyers of type b is sufficiently high relative to the measure of buyers of type c. These conditions are intuitive. The first condition guarantees that there is not enough competition in the product market to drive basket prices below the level $u_b + u_c$. The second condition guarantees that, given α, the number of buyers of type c or the valuation of buyers of type c are low enough to guarantee that a seller does not find it optimal to deviate from one of the basket prices on the support of the distribution $G(q)$, to a pair of prices (p_1, p_2) that induce trade not only with buyers of type b, but also with some buyers of type c.

The analysis of the existence and features of a Bundled Equilibrium is summarized in the following proposition.

Proposition 1: (i) In a Bundled Equilibrium, the marginal distribution of basket prices, G, is continuous over the support $[\underline{q}, \overline{q}]$, with $u_b + u_c < \underline{q} < \overline{q} = 2u_b$, and it is uniquely given by (16). The joint distribution of prices, H, is not uniquely pinned down. (ii) A Bundled Equilibrium exists if and only if condition (18) is satisfied.

Proof: We established part (i) in the main text. We have also established that, given the marginal distribution G in (16), a seller attains a profit of P^* for all $(p_1, p_2) \in R_1$ such that $p_1 + p_2 \in [\underline{q}, \overline{q}]$. In order to complete the proof of part (ii), all we need to do is find a condition under which a seller cannot attain a profit strictly greater than P^* by posting some off-equilibrium prices.

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

$$\begin{aligned} P_1(p_1 + p_2) &= \mu_b \{\alpha + 2(1-\alpha)\} (p_1 + p_2) \\ &< \mu_b \{\alpha + 2(1-\alpha)\} \underline{q} = P^*, \end{aligned}$$

where the first line makes use of the fact that $G(p_1 + p_2) = 0$ and the second line makes use of the fact that $p_1 + p_2 < \underline{q}$. Hence, a seller never finds it optimal to deviate from the equilibrium and post prices $(p_1, p_2) \in R_1$ with $p_1 + p_2 \in (u_b + u_c, \underline{q})$.

In Lemma 3, we proved that a seller never finds it optimal to post prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$ and both p_1 and p_2 greater than u_c and smaller than u_b. If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in [0, u_c]$ and $p_2 \in (u_c, u_b]$, it attains a profit of

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

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

$$\frac{\mu_c}{\mu_b} \leq \frac{3\alpha - 2}{(2-\alpha)u_c/u_b} - 1. \tag{19}$$

Hence, if and only if (19) holds a seller does not find it optimal to deviate from the equilibrium and post prices (p_1,p_2) with $p_1+p_2 \in (2u_c, u_b+u_c]$, $p_1 \in [0,u_c]$ and $p_2 \in (u_c,u_b]$. Similarly, condition (19) guarantees that a seller does not find it optimal to deviate from the equilibrium and post prices (p_1,p_2) with $p_1+p_2 \in (2u_c, u_b+u_c]$, $p_1 \in (u_c,u_b]$ and $p_2 \in [0,u_c]$.

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

$$\begin{aligned} P(p_1,p_2) &= (\mu_b+\mu_c)\left\{\alpha + 2(1-\alpha)\right\}(p_1+p_2) \\ &\leq (\mu_b+\mu_c)\left\{\alpha + 2(1-\alpha)\right\} 2u_c \\ &= P(u_c,u_c), \end{aligned}$$

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

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha}{(2-\alpha)u_c/u_b} - 1. \tag{20}$$

Hence, if and only if (20) holds a seller does not find it optimal to deviate from the equilibrium and post prices (p_1,p_2) with $p_1 \in [0,u_c]$ and $p_2 \in [0,u_c]$. Finally, notice that if condition (19) holds, so does condition (20). Therefore, a seller does not want to deviate from the equilibrium if and only if condition (19), which is the same as condition (18), is satisfied. ■

3.4 Discrimination Equilibrium

In this section, we characterize the properties and derive the conditions for the existence of an equilibrium in which the basket price set by some sellers is strictly greater than u_b+u_c, the basket price set by some other sellers is smaller than u_b+u_c, but all of the sellers set a basket price strictly greater than $2u_c$. That is, we focus on an equilibrium in which $G(u_b+u_c) \in (0,1)$ and $G(2u_c) = 0$. We shall denote this type of equilibrium as a Discrimination Equilibrium, for reasons that will soon be obvious.

We start the characterization of a Discrimination Equilibrium by looking at the distribution of sellers who set a basket price strictly greater than $u_b + u_c$. As the arguments developed in the Section on Bundled Equilibrium also apply to the sellers posting a basket price greater than $u_b + u_c$ in a Discrimination Equilibrium, it follows that, among these sellers, the support of the marginal distribution of basket prices is the closed interval $[q^*, \overline{q}]$, with $q^* < \overline{q} = 2u_b$. Hence, a seller who sets a basket price of $2u_b$ must attain the maximized profit P^*, i.e. $P_1(2u_b) = P^*$ or equivalently

$$P^* = \mu_b \alpha 2u_b. \tag{21}$$

Similarly, every seller who sets a basket price q in the interval $[q^*, \overline{q}]$ must attain the profit P^*, i.e. $P_1(q) = P^*$. This equal profit condition pins down the marginal price distribution of basket prices. In particular, for all $q \in [q^*, \overline{q}]$, $G(q)$ is given by

$$G(q) = 1 - \frac{\alpha}{2(1-\alpha)} \frac{2u_b - q}{q}. \tag{22}$$

Next, we characterize the distribution of those sellers who post a basket price q between $2u_c$ and $u_b + u_c$. Section 3.2 gives us some useful information about the behavior of sellers in this region. Lemma 3 implies that all of these sellers post prices $(p_1, p_2) \in R_2$, i.e. (p_1, p_2) such that the price of one good is smaller than the valuation of buyers of type c, and the price of the other good is strictly greater than the valuation of buyers of type c and smaller than the valuation of buyers of type b. Part (i) of Lemma 4 implies that the marginal distribution of basket prices among these sellers does not have any mass points. Part (ii) of Lemma 4 implies that the marginal distribution of the price of the cheap good among these sellers does not have any mass points (except possibly at zero). From these observations, it follows that the profit enjoyed by a seller who sets a basket price $q \in (2u_c, u_b + u_c]$ and sells good i at the price $p \leq u_c$ is given by

$$P_{2i}(q,p) = \mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q)]\right\} q + \mu_c \left\{\alpha + 2(1-\alpha)[1 - F_i(p)]\right\} p. \tag{23}$$

The next lemma shows that the marginal distribution of prices F_i is the same for good 1 and for good 2, or else sellers would be leaving some profit opportunities unexploited. The lemma immediately implies that the seller's profit function P_{2i} is the same for good 1 and for good 2. For this reason, we shall drop the subscript i from both F_i and P_{2i}.

Lemma 7: In a Discrimination Equilibrium, $F_1(p) = F_2(p)$ for all $p \in [0, u_c]$.

Proof: It is easy to verify that in a Discrimination Equilibrium, $F_i(p)$ does not have a mass point at $p = 0$. Hence, $F_i(p)$ is continuous with $F_i(0) = 0$. Now, suppose that there exist prices p' and p'', with $0 \leq p' < p'' \leq u_c$, for which $F_1(p) > F_2(p)$ for all $p \in (p', p'')$ and $F_1(p') = F_2(p')$. In a Discrimination Equilibrium, any price $p \in [0, u_c]$ for good i is posted by a seller in region R_2, whose profit is given by $P_{2i}(q, p)$. Since $F_1(p) > F_2(p)$ for all $p \in (p', p'')$, the seller attains a strictly higher profit by setting a basket price $q \in (2u_c, u_b + u_c]$ and a price $p \in (p', p'')$ for good 2 than by setting the basket price q and the price p for good 1, i.e. $P_{21}(q, p) < P_{22}(q, p)$ for all $p \in (p', p'')$. In turn, $P_{21}(q, p) < P_{22}(q, p)$ implies that in equilibrium there are no sellers posting a price $p \in (p', p'')$ for good 1, i.e. $F_1(p) = F_1(p')$ for all $p \in (p', p'')$. On the other hand, since F_2 is a distribution function, $F_2(p) \geq F_2(p')$ for all $p \in (p', p'')$. Since $F_2(p') = F_1(p')$, $F_2(p) \geq F_1(p)$ for all $p \in (p', p'')$. We have thus reached a contradiction. ■

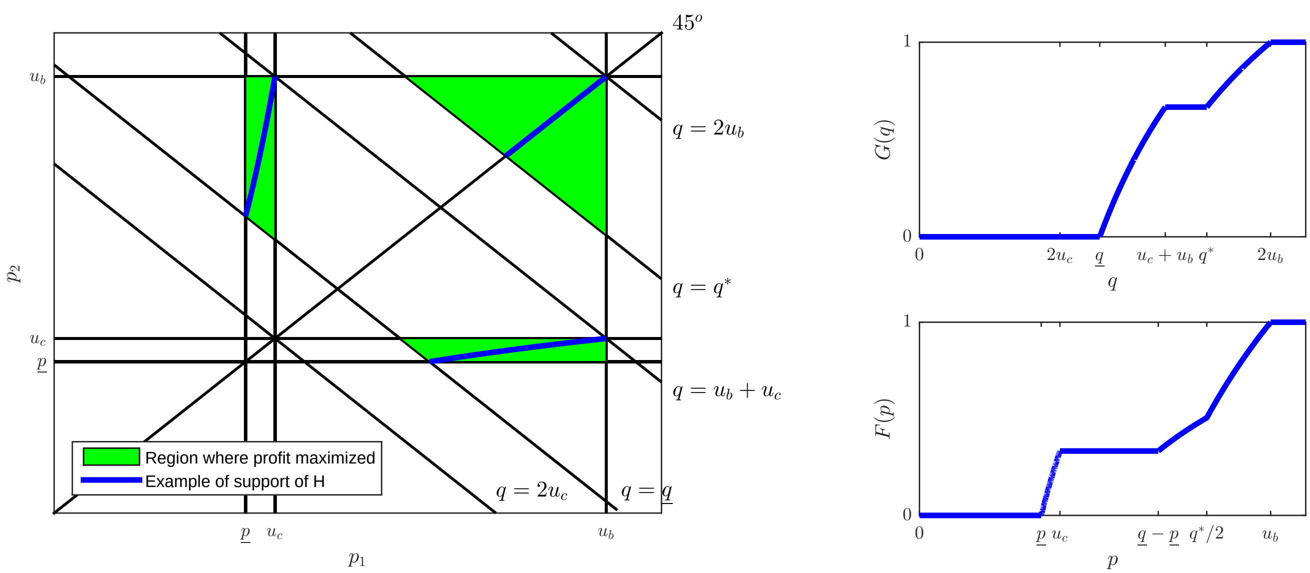

Notes: This figure shows the possible range of the support of the joint distribution H, the shape of the cumulative distributions $G(q)$ and an example of the shape of the cumulative distribution $F(p)$, in the Discrimination Equilibrium.

Figure 5: Discrimination Equilibrium support of $H(p_1, p_2)$ and shape of $G(q), F(p)$

The next lemma is the key to the characterization of a Discrimination Equilibrium. In particular, the lemma establishes that the seller's profit function in (23) is constant for all (q, p) with $q \in [\underline{q}, u_b + u_c]$, $p \in [\underline{p}, u_c]$ and $(p, q - p) \in R_2$, where $\underline{q}$ denotes the lower bound on the support of the marginal distribution of basket prices and $\underline{p}$ denotes the lower bound on the support of the marginal distribution of prices for an individual good. The proof of the lemma is lengthy and relegated into the appendix.

Lemma 8: In a Discrimination Equilibrium, $P_2(q,p) = P_2(u_b + u_c, u_c)$ for all (q,p) with with $q \in [\underline{q}, u_b + u_c]$, $p \in [\underline{p}, u_c]$ and $(p, q-p) \in R_2$.

First, Lemma 8 implies that a seller posting prices (p_1, p_2) with $p_1 \in [\underline{p}, u_c]$, $p_2 \in (u_c, u_b]$ and $p_1 + p_2 \in (\underline{q}, u_b + u_c]$ attains the same profit as a seller posting the price u_c for good 1 and u_b for good 2. In equilibrium, there are some sellers posting prices (p_1, p_2) with $p_1 \in [\underline{p}, u_c]$, $p_2 \in (u_c, u_b]$ and $p_1 + p_2 \in (\underline{q}, u_b + u_c]$ and their profit must be equal to P^*. Taken together, these observations imply that $P_2(u_b + u_c, u_c) = P^*$ or, equivalently,

$$\mu_b \left\{\alpha + 2(1-\alpha)[1 - G(u_b + u_c)]\right\}(u_b + u_c) + \mu_c \left\{\alpha + 2(1-\alpha)[1 - F(u_c)]\right\} u_c = P^*. \quad (24)$$

Similarly, a seller that posts the lowest equilibrium basket price among those greater than $u_b + u_c$ must also attain a profit of P^*. That is, $P_1(q^*) = P^*$ or, equivalently,

$$\mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q^*)]\right\} q^* = P^*. \quad (25)$$

Equating the left-hand sides of (24) and (25), we obtain

$$\begin{aligned} &\mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q^*)]\right\}(u_b + u_c) + \mu_c \left\{\alpha + 2(1-\alpha)[1 - G(q^*)/2]\right\} u_c \\ &= \mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q^*)]\right\} q^*, \end{aligned} \quad (26)$$

where we used the fact that $G(u_b + u_c) = G(q^*)$, as the measure of sellers setting a basket price q between $u_b + u_c$ and q^* is zero, and using the fact that $F(u_c) = G(u_b + u_c)/2$, as the measure of sellers who post a price for good 1 below u_c are one half of the measure of sellers who set a basket price below $u_b + u_c$. The equation in (26) can be solved with respect to q^* to obtain

$$q^* = \frac{2\alpha(1 + u_c/u_b) + \alpha(\mu_c/\mu_b)(u_c/u_b)}{4\alpha - (2-\alpha)(\mu_c/\mu_b)(u_c/u_b)} 2u_b. \quad (27)$$

Second, Lemma 8 implies that $P_2(q, u_c) = P_2(u_b + u_c, u_c)$ for all $q \in [\underline{q}, u_b + u_c]$. That is, for all $q \in [\underline{q}, u_b + u_c]$, we have

$$\begin{aligned} &\mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q)]\right\} q + \mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F(u_c)\right]\right\} u_c \\ &= \mu_b \left\{\alpha + 2(1-\alpha)[1 - G(q^*)]\right\}(u_b + u_c) + \mu_c \left\{\alpha + 2(1-\alpha)\left[1 - F(u_c)\right]\right\} u_c. \end{aligned} \quad (28)$$

We can solve the functional equation in (28) with respect to the marginal distribution of basket prices among sellers in the region R_2. We then find that, for all $q \in [\underline{q}, u_b + u_c]$, $G(q)$ is given by

$$G(q) = G(q^*) - \frac{\alpha + 2(1-\alpha)[1 - G(q^*)]}{2(1-\alpha)} \frac{u_b + u_c - q}{q}. \quad (29)$$

The solution to the equation $G(q) = 0$ is $\underline{q}$, the lower bound of the support of the marginal distribution of basket prices. In particular, $\underline{q}$ is given by

$$\underline{q} = \frac{2\alpha u_b}{2-\alpha}\frac{u_b + u_c}{q^*}. \tag{30}$$

Third, Lemma 8 implies that $P(\underline{q}, p) = P(\underline{q}, u_c)$ for all $p \in [\underline{p}, u_c]$. That is, for all $p \in [\underline{p}, u_c]$, we have

$$\begin{aligned} &\mu_b \{\alpha + 2(1-\alpha)\} \underline{q} + \mu_c \{\alpha + 2(1-\alpha)[1 - F(p)]\} p \\ &= \mu_b \{\alpha + 2(1-\alpha)\} \underline{q} + \mu_c \{\alpha + 2(1-\alpha)[1 - F(u_c)]\} u_c. \end{aligned} \tag{31}$$

The functional equation (31) can be solved with respect to the marginal distribution of prices for an individual good. We then find that, for all $p \in [\underline{p}, u_c]$, $F(p)$ is given by

$$F(p) = G(q^*)/2 - \frac{\alpha + 2(1-\alpha)\left[1 - G(q^*)/2\right]}{2(1-\alpha)}\frac{u_c - p}{p}, \tag{32}$$

where we used the fact that $F(u_c) = G(q^*)/2$. The solution to the equation $F(p) = 0$ is $\underline{p}$, the lower bound on the support of the marginal distribution of prices for an individual good. In particular, $\underline{p}$ is given by

$$\underline{p} = \frac{\alpha + 2(1-\alpha)\left[1 - G(q^*)/2\right]}{2-\alpha} u_c. \tag{33}$$

This completes the characterization of a Discrimination Equilibrium, which is illustrated in Figure 5. In any such equilibrium, the marginal distribution of basket prices, $G(q)$, is uniquely pinned down and it is given by (22) for $q \in [q^*, \overline{q}]$ and by (29) for $q \in [\underline{q}, u_b + u_c]$. Over both intervals, the shape of the distribution is the same as in a version of Burdett and Judd (1983) where the buyer's valuation for the only good traded is $2u_b$. This property of equilibrium is intuitive. Over both intervals, the distribution G keeps the profit that the seller makes from buyers of type b constant. Buyers of type b must purchase both goods together and, hence, they only care about the price of the basket relative to their valuation of the basket $2u_b$. Therefore, the distribution that keeps the profit that a seller makes off of buyers of type b constant has the same shape as in a model where there is a single good worth $2u_b$. However, unlike in Burdett and Judd (1983), G has a gap in the support because $q^* > u_b + u_c$. Also this property of equilibrium is easy to understand. A seller setting a basket price q^* trades with buyers of type b, as for each good it posts a price above the valuation of buyers of type c. A seller setting a basket price equal to $u_b + u_c$ trades with the same number of buyers of type b, however it also trades one good to at least $\mu_c \alpha$ buyers of type c. The two sellers can attain the same total profit only if q^* is strictly greater than $u_b + u_c$.

The marginal distribution of prices for an individual good, $F(p)$, is uniquely pinned down over the interval $[\underline{p}, u_c]$ and it is given by (32), i.e. the interval where the price of an individual good is allocative. Notice that, over this interval, the shape of the distribution is the same as in a version of Burdett and Judd (1983) where the buyer's valuation for the good is u_c. The distribution F keeps the profit that the seller makes off of buyers of type c constant. Buyers of type c only purchase one good, for which they have valuation u_c. Therefore, the shape of the distribution that keeps the profit that the seller makes off of this type of buyers constant is the same as in a model where a single good worth u_c is traded.

The joint price distribution H is not uniquely pinned down. For sellers with $q \in [q^*, \overline{q}]$, any joint distribution H that has support inside R_1 and generates the marginal distribution of basket prices G is consistent with equilibrium. For example, the joint distribution H in which every seller with a $q \in [q^*, \overline{q}]$ posts the price $q/2$ for both good 1 and good 2 is consistent with equilibrium. For sellers with $q \in [\underline{q}, u_b + u_c]$, any joint distribution H that has support inside R_2 and simultaneously generates the marginal distribution of basket prices G and the marginal distribution of prices for individual goods F is consistent with equilibrium. For example, the joint distribution H in which the sellers with $p_1 \in [\underline{p}, u_c]$ post prices $(p_1, \phi(p_1) - p_1)$ and sellers with $p_2 \in [\underline{p}, u_c]$ post prices $(\phi(p_2) - p_2, p_2)$ is consistent with equilibrium, where $\phi(p)$ is the basket price for a seller with a lowest good price of p and it is given by

$$\phi(p) = \frac{[\alpha + 2(1-\alpha)(1-G(q^*))]\,(u_b + u_c)}{[\alpha + 2(1-\alpha)(1-G(q^*))] + 2\,[\alpha + 2(1-\alpha)(1-G(q^*)/2)]\,(u_c - p)/p} \tag{34}$$

In a Discrimination Equilibrium, as in a Bundled Equilibrium, sellers post different prices for the same basket of goods and, in this sense, there is dispersion in store prices. In a Discrimination Equilibrium, unlike in a Bundled Equilibrium, there is always dispersion across sellers in the price of good 1 relative to good 2 (or, equivalently, with respect of the basket price). Indeed, there is always (at least) a measure $G(q^*)/2$ of sellers with a price of good 1 relative to good 2 that is strictly smaller than 1 (i.e. the sellers in R_2 above the 45 degree line), and (at least) a measure $G(q^*)/2$ of sellers with a price of good 1 relative to good 2 that is strictly greater than 1 (i.e. the sellers in R_2 below the 45 degree line). Hence, in any Discrimination Equilibrium, there is always relative price dispersion.

It is easy to explain why relative price dispersion emerges as an equilibrium outcome. Competition between sellers drives the distribution of basket prices down to the region where

q is between $2u_c$ and $u_b + u_c$. In this region, a seller never finds it optimal to post the same price $q/2$ for both goods. Instead, a seller finds it optimal to set the price of one good below the valuation of buyers of type c and the price of the other good above the valuation of buyers of type c. That it, a seller finds it optimal to follow an asymmetric pricing strategy. In equilibrium, if there are some sellers posting a higher price for good 1 than for good 2, then some other sellers must post a higher price for good 2 than for good 1, or else there would be some unexploited profit opportunities. That is, equilibrium requires symmetry in the cross-sectional distribution of prices. The asymmetric pricing strategy followed by each individual sellers combined with the symmetry of the overall distribution of sellers leads, necessarily, to relative price dispersion.

The asymmetric pricing strategy followed by individual sellers is a way to price discriminate between the two types of buyers. The difference between the two types of buyers in the valuation of the goods creates an incentive to price discriminate. Indeed, if the seller could distinguish between busy and cool buyers, he would offer them different prices. The difference between the two types of buyers in the willingness to visit multiple retailers creates the opportunity to price discriminate. Indeed, since the busy buyers have high valuation and need to purchase everything in the same store, while the cool buyers have low valuation and are willing to shop at multiple stores, the seller can charge a low price (on some items) to the cool buyers and a high price (on average) to the busy buyers by pricing its goods asymmetrically.

It is interesting to contrast the type of price discrimination described above with intertemporal price discrimination (see, e.g., Conslik, Gerstner and Sobel 1984 and Sobel 1984 or, in a search-theoretic context, Albrecht, Postel-Vinay and Vroman 2013 and Menzio and Trachter 2015). 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.

In the Discrimination Equilibrium (as in any other equilibrium of the model), a seller is indifferent between posting any prices (p_1, p_2) on the support of the joint distribution H. Hence, in a repeated version of the model, an individual seller might sample a different pair of prices in every period. In this case, every seller would, on average over a sufficiently long interval of time, charge approximately the same price for every good and, therefore, the same relative price for one good compared to the other. Relative price dispersion would vanish. However, there are several reasons to believe that sellers would not find it optimal to frequently redraw their prices.[9] For instance, if sellers face menu costs, they strictly prefer not changing their nominal prices rather than to drawing a new pair from the distribution H in every period. The real prices would move only to the extent that they are eroded by inflation, unless they reach some lower bound that induces the seller to pay the menu cost.[10] Alternatively, if there is some small amount of heterogeneity across sellers with respect to their wholesale costs, the mixed strategy gets "purified." That is, every individual seller has strict preferences over a particular point in the distribution H.

We conclude the analysis of a Discrimination Equilibrium by presenting the necessary and sufficient conditions for its existence. In particular, we can show that this type of equilibrium exists if and only if

$$\frac{\mu_c}{\mu_b} > \frac{3\alpha - 2}{(2-\alpha)u_c/u_b} - 1 \tag{35}$$

and

$$\frac{\mu_c}{\mu_b} \leq \frac{\alpha - (2-\alpha)u_c/u_b}{1 + (2-\alpha)u_c/u_b} \frac{1 + u_c/u_b}{u_c/u_b}. \tag{36}$$

Intuitively, condition (35) guarantees that some sellers find it optimal to post a basket price q lower than $u_b + u_c$. The condition is satisfied when: (i) the retail market is sufficiently competitive, in the sense that α is smaller than 2/3, or when (ii) the measure and/or the valuation of buyers of type c relative to buyers of type b are sufficiently high. Condition (36) guarantees that no seller finds it optimal to set prices (p_1, p_2) with $p_1 + p_2 \leq 2u_c$. The condition is satisfied when: (i) the retail market is not too competitive, in the sense that α is greater than $2(u_c/u_b)/(1 + 2u_c/u_b)$, and (ii) the measure and/or the valuation of buyers of type c relative to the measure of buyers of type b are sufficiently low.

[9] As we documented in Section 2, sellers do change their prices at high frequency. However, these transitory price changes may be related to sales and other forms of systematic intertemporal price discrimination, rather than to new draws from the equilibrium price distribution. Indeed, Eichenbaum, Jaimovich, and Rebelo (2011) document that the behavior of the price of an individual good at a particular seller alternates between a small, relatively stable number of price points, e.g. a baseline price and a sale price.

[10] Burdett and Menzio (2014) present a version of Burdett and Judd (1983) where sellers post nominal prices and face menu costs.

A couple of remarks about the necessary and sufficient conditions for the existence of a Discrimination Equilibrium are in order. First, notice that there is a non-empty region of parameter values for which a Discrimination Equilibrium exists. In fact, the right-hand side of (36) is strictly greater than the right-hand side of (35) and, hence, there is a non-empty region for which both conditions are satisfied. Second, notice that the region where a Discrimination Equilibrium exists and the region where a Bundled Equilibrium exist are non-overlapping and contiguous. In fact, condition (35) is the opposite of (18), the necessary and sufficient condition for a Bundled Equilibrium. Figure 6 illustrates the existence of the two types of equilibria in the space of parameters $\{u_c/u_b, \mu_c/\mu_b\}$.

Figure 6: Equilibrium depending on valuation and measure of buyer types

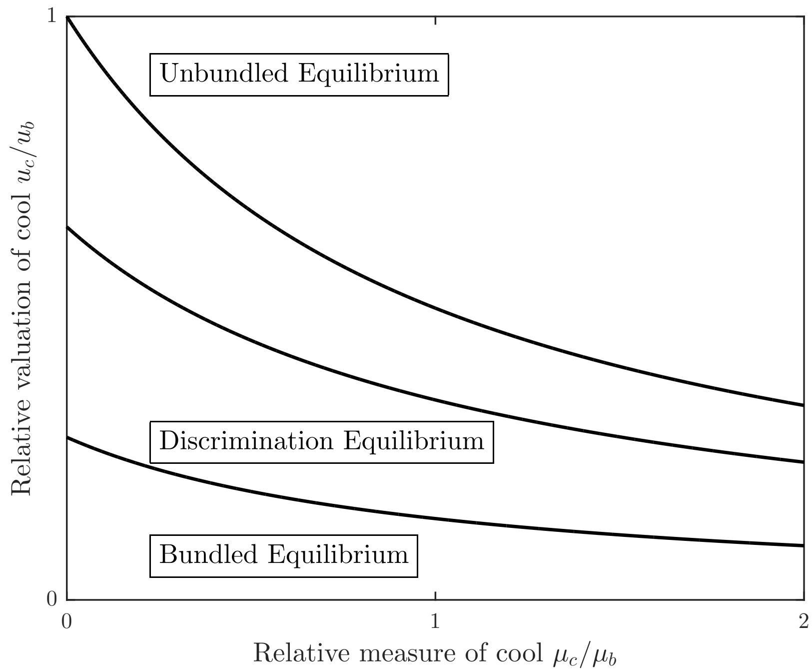

Notes: This figure shows the type of equilibrium occurring depending on the relative valuation and measure of cool buyers in the market.

The analysis of the Discrimination Equilibrium is summarized in the following proposition.

Proposition 2: (i) In a Discrimination Equilibrium, the marginal distribution of basket prices, G, has support $[\underline{q}, u_b+u_c] \cup [q^*, \overline{q}]$, with $2u_c < \underline{q} < u_b+u_c$ and $u_b+u_c < q^* < \overline{q} = 2u_b$; over the interval $[\underline{q}, u_b+u_c]$, G is continuous and it is given by (22); over the interval $[q^*, \overline{q}]$, G is continuous and it is given by (29). The marginal distribution of prices for an individual good, F, has support $[\underline{p}, u_c]$ for $p \in [0, u_c]$. Over the interval $[\underline{p}, u_c]$, F is continuous and it is given by (32). The joint distribution of prices, H, is not uniquely pinned down. (ii) A Discrimination Equilibrium exists if and only if conditions (35) and (36) are satisfied.

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

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

First, we compute the seller's profit for prices (p_1, p_2) with $p_1 + p_2 \in (u_b + u_c, 2u_b]$. If the seller posts prices (p_1+p_2) with $p_1+p_2 \in [q^*, \overline{q}]$, it attains a profit of P^*, as guaranteed by the construction of G and Lemma 2. If the seller posts prices (p_1, p_2) with $p_1+p_2 \in (u_b+u_c, q^*)$, it attains a profit of

$$\begin{aligned} P_1(p_1 + p_2) &= \mu_b \{\alpha + 2(1-\alpha)[1 - G(p_1 + p_2)]\} (p_1 + p_2) \\ &< \mu_b \{\alpha + 2(1-\alpha)[1 - G(q^*)]\} q^* \\ &= P^*, \end{aligned}$$

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

Second, we compute the seller's profit for prices (p_1, p_2) such that $p_1 + p_2 \in (2u_c, u_b + u_c]$. In Lemma 3, we showed that the seller never finds it optimal to post prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (\underline{q}, u_b + u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in [\underline{p}, u_c]$, it attains a profit of P^*, as guaranteed by Lemma 8. Similarly, if the seller posts prices (p_1, p_2) with $p_1+p_2 \in (\underline{q}, u_b+u_c]$, $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$, it attains a profit of P^*. If the seller posts prices (p_1, p_2) with $p_1 + p_2 \in (2u_c, \underline{q})$, $p_1 \in [\underline{p}, u_b]$ and $p_2 \in (u_c, u_b]$, it attains a profit of

$$\begin{aligned} P_2(p_1 + p_2, p_1) &= \mu_b \{\alpha + 2(1-\alpha)\} (p_1 + p_2) + \mu_c \{\alpha + 2(1-\alpha)[1 - F(p_1)]\} p_1 \\ &\leq \mu_b \{\alpha + 2(1-\alpha)\} \underline{q} + \mu_c \{\alpha + 2(1-\alpha)[1 - F(p_1)]\} p_1 \\ &= P^*. \end{aligned}$$

The first line uses the fact that the seller trades both goods with all the buyers of type b it meets. The second line uses the fact that the seller would also trade both goods with all the buyers of type b it meets at the basket price $\underline{q}$, and the last line uses Lemma 8. Similarly, $P_2(p_1+p_2,p_2) \leq P^*$ for all (p_1,p_2) with $p_1+p_2 \in (2u_c, \underline{q})$, $p_1 \in (u_c,u_b]$, $p_2 \in [\underline{p},u_c]$.

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

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

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

If and only if $P(u_c,u_c) \leq P^*$, the highest profit that the seller can attain is P^*. Using the fact that $P_2(u_b+u_c,u_c) = P^*$, we can write the condition $P(u_c,u_c) \leq P^*$ as

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

After substituting out $G(q^*)$, we can write the inequality above as (36).

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

$$\frac{\mu_c}{\mu_b} > \frac{3\alpha-2}{(2-\alpha)u_c/u_b} - 1, \tag{37}$$

and

$$\frac{\mu_c}{\mu_b} < \frac{\alpha(1-u_c/u_b)}{u_c/u_b}. \tag{38}$$

Condition (37) is necessary and sufficient for $G(q^*) > 0$, and it is condition (35). Condition (38) is necessary and sufficient for $q^* < 2u_b$, and it holds whenever condition (36) is satisfied. If and only if (37) and (38) are satisfied, G and F are proper distribution functions. That is, the interval $[\underline{q}, u_b+u_c]$ is non-empty and, over this interval, $G(q)$ is strictly increasing in q, and such that $G(\underline{q}) = 0$ and $G(u_b+u_c) = G(q^*)$, where $G(q^*) \in (0,1)$. The interval $[q^*, \overline{q}]$ is non-empty and, over this interval, $G(q)$ is strictly increasing in q, and such that $G(q^*) = G(u_b+u_c)$ and $G(\overline{q}) = 1$. Similarly, the interval $[\underline{p},u_c]$ is non-empty and, over this interval, $F(p)$ is strictly increasing in p and such that $F(\underline{p}) = 0$ and $F(u_c) = G(q^*)/2 \in (0,1)$.

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

3.5 Other equilibria

Figure 6 displays in the $\{u_c/u_b, \mu_c/\mu_b\}$ space, the area where a Bundled Equilibrium and the area where a Discrimination Equilibrium exist. Fix the valuation of type-c buyers relative to the valuation of type-b buyers, u_c/u_b. If the measure of type-c buyers relative to type-b buyers, μ_c/μ_b, is sufficiently low, there is a Bundled Equilibrium. In this type of equilibrium, every seller sets prices (p_1, p_2) with $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$ and, hence, it might trade the basket of goods to buyers of type b, but it might never trade with buyers of type c. If we increase μ_c/μ_b, we enter the region where there is a Discrimination Equilibrium. In this type of equilibrium, there are some expensive sellers that set prices (p_1, p_2) with $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. These sellers might trade the basket of goods to buyers of type b, but they never trade anything to buyers of type c. Moreover, there are some mid-range sellers that set prices (p_1, p_2) with one price in $[0, u_c]$ and the other price in $(u_c, u_b]$. These sellers might trade the basket of goods to buyers of type b, and the cheaper good to buyers of type c. If we increase μ_c/μ_b further, we find two other types of equilibria. For the sake of brevity, we shall only describe the main features of these equilibria.

If μ_c/μ_b is sufficiently high, there is an Unbundled Equilibrium. In this type of equilibrium, every seller posts prices (p_1, p_2) with $p_1 \in [0, u_c]$ and $p_2 \in [0, u_c]$. Hence, every seller might trade the basket of goods to buyers of type b, or it might trade either one or both goods to buyers of type c. The marginal distribution of prices for each individual good is uniquely pinned down and it has the same shape as in a one-good version of Burdett and Judd where the valuation of each good is u_c. The equilibrium may or may not feature relative price dispersion depending on the choice of the joint distribution of H, which is not pinned down uniquely.

If μ_c/μ_b is below the existence region for an Unbundled Equilibrium and above the existence region for a Discrimination Equilibrium, there is a type of equilibrium in which some expensive sellers post prices (p_1, p_2) with $p_1 \in (u_c, u_b]$ and $p_2 \in (u_c, u_b]$. These sellers might only trade the basket of goods to buyers of type b. There are some mid-range sellers that post prices (p_1, p_2) with one price in $[0, u_c]$ and the other price in $(u_c, u_b]$. These sellers might trade the basket of goods to buyers of type b, and the cheaper good to buyers of type c. Finally, there are some cheap sellers that post prices (p_1, p_2) with $p_1 \in [0, u_c]$ and $p_2 \in [0, u_c]$. These sellers might might trade the basket of goods to buyers of type b, or it might trade either one or both goods to buyers of type c. In this type of equilibrium, there is always relative price dispersion.

4 Validation

In this section, we want to provide some evidence in support of our theory of relative price dispersion. First, our theory implies that different households are going to pay a different amount for purchasing the same basket of goods, in the same period of time and in the same geographical area. Indeed, buyers of a given type purchase from sellers that post different prices and, hence, they end up paying different prices for the same basket of goods. Moreover, different types of buyers are heterogeneous in their ability to exploit relative price dispersion and, hence, they end up paying different prices for the same basket of goods. Second, our theory implies that the households who have a relatively low value of time (i.e. the cool buyers) should pay sistematically lower prices for the same basket for the same basket of goods than the households who have a relatively high value of time (i.e. the busy buyers). Third, our theory implies that the households who visit more store when shopping should pay sistematically lower prices for the same basket for the same basket of goods than the households who have a relatively high value of time (i.e. the busy buyers). We now discuss the evidence in support of each of these implications of our model.

4.1 Dispersion in household price indexes

Kaplan and Menzio (2014a) use household-level scanner data from the Kilts-Nielsen Consumption Panel dataset to document the existence of a large amount of dispersion in the prices that different households paid for identical bundles of goods in the same geographical market. Following closely the methodology in Aguiar and Hurst (2007), they construct household price indexes by comparing each household's quarterly expenditure with their counterfactual expenditure had they purchased each of their chosen goods at the market-wide average price for that good. Table 6 of Kaplan and Menzio (2014a) reports that the standard deviation of household-level price indexes is around 9%, more than half the size of the standard deviation in individual prices. Further, a household at the 90th percentile of the distribution of price indexes pays 22% more on average for the same basket of goods as a household at the 10th percentile of the distribution. Therefore, as predicted by our model, different households do pay different prices for the same basket of goods, in the same period of time and in the same geographical area.

According to our theory, some of the dispersion in the price index of different households is due to dispersion in the average price of different stores. However, according to our theory, there should also be some dispersion in the price index of different households that is due to relative price dispersion. Kaplan and Menzio (2014a), Table 7 shows that slightly less than half of the variance of price indexes can be attributed to differences in the store component of prices. That is, roughly half of the dispersion in the price indexes is due to the fact that some households shop at expensive stores and some shop at cheap stores. The other half of the variance of price indexes can be attributed to differences in the quarterly averages of the store-good component of prices. That is, roughly half of the dispersion in the price indexes is due to the fact that some households are systematically better than others at buying the goods that are cheaper at each of the stores they visit.

4.2 Value of time and prices

According to our theory, there should be systematic differences in the price index of households with a different value of time. Indeed, the buyers with a low value of time (i.e. the cool buyers) should pay lower prices for the same basket of goods than the buyers with a high value of time (i.e. the busy buyers).

Aguiar and Hurst (2007) estimate the relationship between household price indexes and the age of the household, focusing on the difference between retired and working-age households. In this setting, age is a natural proxy for the value of time since retired households spend less time in the labor market than working-age households. They find very little age variation in price indexes for households aged 50 or below, and a sharp decline of price indexes with age for households above age 50. Their estimates suggest that retirement-age households pay around 2% less than working age households for identical bundles of goods. Kaplan and Menzio (2014a) substantiate these findings using a much larger and more representative dataset. Their Figure 9 shows a similar magnitude difference in prices paid between retirement-age and working-age households. Another natural proxy for the value of time is employment status. Kaplan and Menzio (2014a) also regress household price indexes on the employment status of the household head(s). They find that households whose heads are

Table 6: Regression of household price indexes on indicators of multi-stop shopping

	Level	Level	Log	Log
N. Store / Dol.	-0.09099**	-0.10365**	-0.01124**	-0.01291**
	(0.00363)	(0.00349)	(0.00041)	(0.00041)
FE	No	Yes	No	Yes
R^2	0.02041	0.01523	0.01776	0.0074

	Level	Level	Log	Log
N. Store	-0.01196**	-0.00367**	-0.03424**	-0.01577**
	(0.00027)	(0.00018)	(0.00072)	(0.00051)
Expend.	0.00008**	0.00007**	0.00964**	0.01261**
	(0.00000)	(0.00000)	(0.00041)	(0.00044)
FE	No	Yes	No	Yes
R^2	0.02707	0.00235	0.02875	0.00753

	Level	Level	Log	Log
N. Store > 1	-0.02533**	-0.00739**	-0.02948**	-0.01318**
	(0.00078)	(0.00062)	(0.00078)	(0.00060)
Expend.	0.00005**	0.00006**	0.00723**	0.01149**
	(0.00000)	(0.00000)	(0.00041)	(0.00043)
FE	No	Yes	No	Yes
R^2	0.01751	0.00187	0.01933	0.00643

Notes: This table presents results for regressions of household price indexes on indicators of multi-stop shopping: average number of different stores visited per dollar spent in the quarter, number of different stores visited per quarter (conditioning on dollars spent per quarter), and an indicator for visiting multiple stores (conditioning on dollars spent per quarter). Level models have expenditures in levels, and Log models have expenditures in logs. In all regressions: N=880104, clusters=78758.

unemployed pay around 1% less for identical bundles of goods than similar households whose heads are employed.[11]

4.3 One-stop and multi-stop shopping

Our model not only suggests that households with a lower value of time will pay lower prices on average, but it also suggest the mechanism through which these households achieve these

[11] The KNCP data that Kaplan and Menzio (2014a) use and that we use in Section 4.3 only provides approximately annual updates to household demographic and labor market information. This limits the ability to exploit the panel dimension of the data when estimating the effect of employment on prices paid. In contrast, the regressions in Section 4.3 relating number of stores visited to prices paid use data that is available at a daily frequency so are able to exploit the panel dimension of the data set for identification.

prices. In our model the households who pay the low bundle prices are those who visit multiple stores, and purchase each good from the store where it is cheapest. If this mechanism is indeed an important feature of product markets, then we should see that households who shop at a larger number of stores systematically pay lower prices. Moreover, for a given household we should see that the household pays lower prices in time periods that it shops at a larger number of stores. We now turn to household-level scanner data to test these predictions of the model?

Our data comes from the Kilts-Nielsen Consumer Panel dataset. We use the same sample as in Kaplan and Menzio (2014a), and we refer the reader to that paper for a detailed description of the dataset, and details of construction of household-level price indexes. In all specifications, our dependent variable is the household price indexes in a given quarter. Our independent variables are constructed from the number of stores from which a household recorded purchases in a given quarter. We consider various alternatives.

We start by regressing household price indexes on the level and log of the number of stores per dollar of expenditure. In all specifications we control for household size, the age and education of household members and a full set of quarter and market dummies. The first column of Table 6 shows that households who visit a larger number of stores per dollar spent pay significantly lower prices for the same bundles. Visiting one additional store per dollar spent is associated with a 9% reduction in prices paid. In the second column of Table 6 we control for household-specific fixed effects. The effect is slightly larger and also strongly significant. Thus the relationship between number of stores and price indexes is not due to unobserved differences in households. Rather, these results imply that the same households pay lower prices for the same goods in periods when they do their shopping in a larger number of stores. In the final two columns of Table 6, we construct the independent variable as the natural logarithm of number of stores per dollar spent. The relationship is strong and significant also according to this measure. Controlling for household fixed effects, the estimates imply that a doubling of the number of stores visited per dollar is associated with around a 1% reduction in prices paid.

One may worry that using number of stores per dollars as the independent variable may impose to strict a relationship between stores and quantities. An alternative approach is to regress quarterly price indexes on the number of stores visited by the household in that quarter, while controlling for the quarterly expenditure of the household. We report results for this alternative specification in the middle panel of Table 6, both in logs and levels, and with and without household fixed effects. The effect of number of stores visited is strong and statistically significant also by this measure.

In fact our model relies on an even stronger feature of shopping behavior. The key distinction is one-stop vs multi-stop shoppers. Motivated by this, in the bottom panel of Table 6, we regress household price indexes on an indicator variable for whether a household shopped at one or more than one stores in a given quarter. The results are again strong and significant. Exploiting both cross-sectional and panel variation, the estimates suggest that households with the same total expenditure who visited more than one store pay around 2.5% less. Exploiting only within-household variation over time, the estimates suggest that a household pays around 1% less in a quarter when it visits more than one store than in a quarter when it visits only one store.

5 Conclusions

TBW

References

Aguiar, M., and E. Hurst (2007): "Lifecycle Prices and Production," *American Economic Review*, 97(5), 1533–59.

Albrecht, J., F. Postel-Vinay, and S. Vroman (2013): "An Equilibrium Search Model of Synchronized Sales," *International Economic Review*, 54(2), 473–93.

Baughman, G., and K. Burdett (2015): "Multi-Product Search Equilibria," Mimeo, University of Pennsylvania.

Blundell, R., and I. Preston (1998): "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113(2), 603–640.

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

Burdett, K., and G. Menzio (2014): "The (Q,S,s) Pricing Rule," NBER Working Paper 19094.

Butters, G. R. (1977): "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 44(3), 465–491.

Conlisk, J., E. Gerstner, and J. Sobel (1984): "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, 99(3), 489–505.

Eden, B. (2013): "Price dispersion and demand uncertainty: Evidence from us scanner data," Mimeo, Vanderbilt University.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011): "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 101, 234–262.

Ellison, G. (2005): "A Model of Add-on Pricing," *Quarterly Journal of Economics*, 120(2), 585–637.

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.

Hotelling, H. (1929): "Stability in Competition," *Economic Journal*, 39, 41–57.

Kaplan, G., and G. Menzio (2014a): "The Morphology of Price Dispersion," *International Economic Review, forthcoming.*

——— (2014b): "Shopping Externalities and Self-Fullfilling Unemployment Fluctuations," *Journal of Political Economy, forthcoming.*

Lal, R., and C. Matutes (1989): "Price Competition in Multimarket Duopolies," *The RAND Journal of Economics*, 20(4), 516–537.

——— (1994): "Retail Pricing and Advertising Strategies," *Journal of Business*, 67(3), 345–370.

McAfee, R. P. (1995): "Multiproduct Equilibrium Price Dispersion," *Journal of Economic Theory*, 67, 83–105.

Menzio, G., and N. Trachter (2015): "Equilibrium Price Dispersion Across and Within Stores," PIER Working Paper 15-003.

Pratt, J. W., D. A. Wise, and R. Zeckhauser (1979): "Price Differences in Almost Competitive Markets," *Quarterly Journal of Economics*, 93(2), 189–211.

Rhodes, A. (2015): "Multiproduct Retailing," *Review of Economic Studies*, 82(1), 333–359.

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.

Stigler, G. J. (1961): "The Economics of Information," *Journal of Political Economy*, 69(3), 213–225.

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

Zhou, J. (2014): "Multiproduct Search and the Joint Search Effect," *American Economic Review*, 104(9), 2918–2939.

Appendix

A Empirical appendix

Table 7: UPCs

Department	Percentage of UPCs (%)
Dairy	21.9
Deli	1.8
Dry Grocery	52.2
Fresh Produce	6
Frozen Foods	8.5
General Merchandise	0.9
Health and Beauty	0.4
Non-Food Grocery	4.3
Packaged Meat	4

Notes: The number and percentage of UPCs in the departments with the highest representation among these 1000 UPCs.

Table 8: UPCs

Product group	Percentage of UPCs (%)
Yogurt	10.7
Carbonated Beverages	9.3
Fresh Produce	6
Bread and Baked Goods	5.3
Pizza / Snacks / Hors doerves - Frozen	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

Notes: The number and percentage of UPCs in the 20 product groups with the highest representation among these 1000 UPCs.