

Discriminating Dispersion in Prices^{*}

G. Kaplan G. Menzio L. Rudanko N. Trachter

March 25, 2015

Preliminary

Abstract

We extend the literature on price dispersion both empirically and theoretically: First, we use a comprehensive dataset on US retail pricing to document that there exists significant dispersion in the overall price level across stores, as well as in the prices of individual goods relative to the store average. While the overall price-level varies by approximately 6% across stores, the prices of individual goods vary by as much as 14% around the store average, making the latter form of 'relative price dispersion' dominant. This suggests non-trivial scope for households to economize by visiting multiple stores and buying goods where they are cheapest. To rationalize these observations, we propose a novel theory of equilibrium multi-product search. Relative price dispersion arises in the model as a result of spatial price discrimination by stores seeking to take advantage in heterogeneity households' willingness to shop around for individual goods.

JEL classification: L11, D40, D83, E31

Keywords: frictional product markets, search, price dispersion, consumption

1 Introduction

There is dispersion in the prices of individual goods across sellers. This observation – even if not backed up by broad empirical evidence at the time – has given rise to an entire literature seeking to generate price dispersion as an equilibrium outcome in a market with search frictions.¹ In this paper we extend that literature both theoretically and empirically: First,

^{*}Kaplan: Princeton University. Email: gkaplan@princeton.edu. Menzio: University of Pennsylvania. Email: gmenzio@econ.upenn.edu. Rudanko: Federal Reserve Bank of Philadelphia. Email: leena.rudanko@gmail.com. Trachter: Federal Reserve Bank of Richmond. Email: nicholas.trachter@rich.frb.org. We are grateful to our audience at the Federal Reserve Bank of Minneapolis 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.

¹E.g., Salop and Stiglitz (1976), Reinganum (1979), Wilde and Schwartz (1979), Burdett and Judd (1983).

we make use of the comprehensive micro-level retail price data available today to quantify and characterize the dispersion in prices households face. Second, motivated by this evidence where purchases typically take the form of baskets of goods rather than individual goods alone, we propose a new theory of price dispersion in a setting with equilibrium multi-product search. Our theory is designed to capture two robust features of the data: i) dispersion in the overall price-level across stores, as well as ii) dispersion in the relative prices of individual goods across stores of a similar overall price level.

Our data come from the Kilts-Nielsen Retail Scanner Dataset, which provides weekly price information for 2.6 million UPCs for the period 2006-11, covering more than half the sales of US grocery and drug stores nationwide. Using this data, we decompose differences in individual goods' prices into a component measuring the overall price-level of the store and a component measuring the price of the individual good relative to the overall price-level of the store. Quantifying the degree of dispersion in the two components separately, we find that the price-levels of stores vary by approximately 6 percent across stores, while individual goods' prices vary by approximately 14 percent around the store price level (at the UPC level). In other words, while both forms of price dispersion are substantial in magnitude, the latter dominates. Stores that are more expensive are not uniformly more expensive across goods, but there is significant variation in how stores price individual goods.

These differences in prices pose a puzzle: there would appear to exist non-trivial scope for households to economize by shopping in the right stores, and especially by shopping around for individual goods to buy them where they are cheapest. The evidence calls for a theory to rationalize the observations.

To this end, we propose a theory of equilibrium multi-product search, where a continuum of stores sell two goods to a continuum of households. We characterise outcomes focusing on two equilibria which illustrate how multi-product search differs from one-product search. The first provides a starting point showing what happens if households shop for baskets instead of individual goods, but must purchase the entire basket from a single store. In this case the equilibrium pins down a distribution of basket prices, but because everyone purchases baskets, the distribution of individual goods' prices is not pinned down. The

second equilibrium goes a step further by allowing some households to purchase the individual items from multiple stores, while others continue to purchase their entire basket in a single store. In this case the equilibrium pins down also the distribution of individual goods' prices (in the range of prices where individual goods are sold). Both equilibria feature dispersion in basket prices, whereas while the first allows for dispersion in relative prices, the second implies it.

In a setting where households shop for, and stores sell, baskets of goods, the basket price may be interpreted as corresponding to the overall price level of the store, and dispersion in this overall price-level thus be viewed as corresponding to the standard notion of price dispersion in a one-dimensional setting. The second form of price dispersion we document – dispersion in the relative prices of individual goods – requires taking the multi-good nature of the problem more seriously. Our theory of relative price dispersion is based on a notion of “spatial price discrimination,” where households are heterogeneous in their opportunity cost of time: some have a relatively high willingness to pay for goods, but are not able to shop around for the individual goods in their basket, while others have a low willingness to pay for goods, and can shop around. This heterogeneity gives stores incentives to play around with their pricing to generate dispersion in relative prices, as some stores price specific goods low in order to attract those shoppers who are not willing to pay as much, but shop around for individual goods.

One concern with the estimates above is that temporary sales could potentially generate a non-trivial degree of dispersion in weekly price data. If the bulk of the dispersion were due to sales, it would be more appropriate to focus on theories of intertemporal price discrimination to explain the evidence, than the static spatial price discrimination theory we propose. To shed light on the importance of such transitory changes in explaining price differences, we adopt a novel approach to studying price data, borrowing from literature studying income and earnings processes (Gottschalk and Moffitt 1994, Blundell and Preston 1998): We decompose changes in both the store price-level and the prices of individual goods relative to the store price-level into persistent and transitory components, implementing a GMM estimation to determine the relative importance of each. We find that the dispersion in store

price levels is essentially entirely persistent, with 94 percent of changes in store price levels persistent. The relative prices of individual goods are more affected by transitory shocks, but with a substantial 37 percent of changes persistent in this case as well. In sum, taking into account the potential role of sales and other short run price changes does not change the conclusion that a substantial degree of price dispersion remains.

Our paper is organized as follows. Section 2 discusses the related literature. Section 3 turns to the data on retail pricing to document the evidence on relative price dispersion and consumer shopping behavior. Section 4 presents our model. Section ?? parameterizes the model and compares to data. Section 5 concludes.

2 Related literature

The existing evidence on price dispersion consists of a set of earlier studies focusing on a narrower set of goods, such as Stigler (1961) and Pratt, Wise, and Zeckhauser (1979), together with two papers making use of the broader retail pricing data available more recently (Eden 2013, Kaplan and Menzio 2014). Most closely related, Kaplan and Menzio (2014) study the Kilts-Nielsen consumer panel data, documenting properties of the distributions of individual goods prices: the shape of the distribution and degree of dispersion by and across goods. We make use of the same data when looking at patterns in consumer spending, but our main focus is on the significantly broader Kilts-Nielsen retail scanner data, with information on listed prices directly from stores, as opposed to purchased prices reported by consumers. We also focus on the relative price dispersion phenomenon, in particular, and to sharpen the measurement develop an approach for quantifying the share of dispersion that is persistent in nature. To this end, our empirical approach makes use of an approach developed for studying the properties of income and earnings processes, to quantify the share of transitory versus persistent variation, in Gottschalk and Moffitt (1994), Blundell and Preston (1998), and subsequent work.

We contribute also by proposing a novel theory of multi-product search, designed to capture the idea of stores taking advantage of consumers' ability to shop around. The literature on

product market search largely focuses on the search for a single good, with only a handful of exceptions. There exists a set of papers modeling the one-sided optimal stopping problem of a consumer searching for multiple products (Burdett and Malueg 1981, Carlson and McAfee 1984, Gatti 1999). These papers are mainly concerned with whether optimal search behavior can be characterized by some type of reservation price policy in this context. McAfee (1995) brings the problem into an equilibrium setting by considering a multi-product extension of Burdett and Judd (1983), and shows that a number of equilibria are possible, and qualitatively different from the unique equilibrium in the single good case. Zhou (2014) offers a recent contribution arguing that pricing decisions in a multi-product search market are affected by a “joint search effect”: with economies of scale in search, setting low prices to attract customers becomes more profitable, as those customers may end up purchasing an entire basket from the store, instead of shopping around. Rhodes (2015) studies the price setting behavior of a monopolist selling multiple goods in a frictional market, emphasizing the hold-up problem consumers are subject to when visiting a multi-product seller based on the advertised prices of particular goods. These papers are theoretical in nature, whereas we have sought to provide evidence as well.

3 Evidence

We begin by explaining our empirical approach – our decomposition of prices dispersion into a store and store-good component, and further into persistent and transitory variation – as well as describing the data, before discussing the results.

3.1 Framework and Estimation

Let p_{jst} denote the quantity-weighted average price of a good j at a store s in time period t . In our application a time period is one week and goods are defined at the UPC (barcode)

level. We assume that the price p_{jst} can be expressed as:

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

$$y_{st} = y_s^F + y_{st}^P + y_{st}^T \quad (2)$$

$$z_{jst} = z_{js}^F + z_{jst}^P + z_{jst}^T \quad (3)$$

The natural logarithm of the quantity-weighted average price of good j at a store s in time period t is comprised of three additively separable components: a component that reflects that average price of the good (and so captures any differences in the units in which a good is measured), a component that reflects the expensiveness of the store selling the good, and a component that reflects factors that are unique to the combination of store and good. The store component and store-good components are assumed to each be comprised of fixed, persistent and transitory components. We assume that the persistent components follow AR(1) process, and that the transitory components follow MA(q) processes. The MA components are necessary because (i) temporary sales may last longer than 1 week, and (ii) the timing of sales may overlap across weekly observation periods.

$$y_{st}^P = \rho_y y_{s,t-1}^P + \eta_{st}^y \quad (4)$$

$$z_{st}^P = \rho_z z_{js,t-1}^P + \eta_{jst}^z \quad (5)$$

$$y_{st}^T = \epsilon_{st}^y + \sum_{i=1}^q \theta_i \epsilon_{s,t-i}^y \quad (6)$$

$$z_{st}^T = \epsilon_{st}^z + \sum_{i=1}^q \phi_i \epsilon_{s,t-i}^z \quad (7)$$

$$y_s^F = \alpha_s^y \quad (8)$$

$$z_{js}^F = \alpha_s^z \quad (9)$$

In our baseline model, we assume that the shocks, $(\alpha^y, \alpha^z, \eta^y, \eta^z, \epsilon^y, \epsilon^z)$ are mean zero white noise random variables that are independent across goods, stores and time.²

²In the appendix, we also consider an alternate specification where we explicitly allow for the fact that the temporary store-good component is likely to reflect sales by modeling it as a 2-point distribution.

We consider estimation of this model given data on quantity-weighted average prices, p_{jst} , for a large number of goods $j = 1 \dots J$, at a large number of store $s = 1 \dots S$ in a market m at a weekly frequency $t = 1 \dots T$. Given the large number of goods, stores and time periods, estimating this model via maximum likelihood is infeasible. Panel data IV regressions are also infeasible because of the unobserved components in prices. Instead we propose to estimate the model using a multi-stage GMM approach analogously to that used in the wage dynamics literature. The parameters to be estimated are the variances of the shocks $(\sigma_{\eta^y}^2, \sigma_{\eta^z}^2, \sigma_{\varepsilon^y}^2, \sigma_{\varepsilon^z}^2)$, the AR(1) parameters (ρ_y, ρ_z) , the MA(q) parameters (θ, ϕ) and the good-time means (μ_{jt}) . We estimate the model in four steps.

In the first step we estimate the good-time means, μ_{jt} , as the market-wide average (over stores, equally weighted) of the quantity-weighted average prices for goods at stores:

$$\mu_{jt} = \frac{1}{S} \sum_{s=1}^S \log p_{jst} \quad (10)$$

We then form demeaned log prices as

$$\hat{p}_{jst} = \log p_{jst} - \hat{\mu}_{jt} \quad (11)$$

In the second step we estimate the store components by taking sample means across all goods in the store

$$\hat{y}_{st} = \frac{1}{n_{jst}} \sum_{j=1}^{n_{js}} \hat{p}_{jst} \quad (12)$$

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

$$\hat{z}_{jst} = \hat{p}_{jst} - \hat{y}_{st} \quad (13)$$

This results in 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)).

In the third step we construct the auto-covariance matrix of each of these panels up to L lags.³

In the fourth step we minimize the distance between the theoretical auto covariance matrices implied by the model and the empirical auto-covariance function from step three. To obtain standard errors we plan on bootstrapping the whole procedure because we are using estimated means as data in the minimization step.

3.2 Data

Our data comes from 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 2011. In order to keep the size of the analysis manageable, we restrict attention to only 1000 commonly sold goods and we estimate the model separately by state. The 1000 UPCs were chosen as the 1000 UPCs with the largest quantities of sales in the Minneapolis designated market area in the first quarter of 2010. These 1,000 products span 50 product groups. For comparison, there was an overall total 189,412 unique UPCs sold in the Minneapolis market in the KNRS data in this quarter. Table 1 shows how these 1000 sample UPCs are distributed across departments.

Table 2 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 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.

For each state, we select stores, goods and weeks that satisfy two criteria: i) For a each

³To allow for time variation in the parameters of the model, we also compute the auto-covariance functions separately for each time periods.

Table 1: Data by department

	Number of UPCs	Percentage of UPCs (%)
Dairy	219	21.9
Deli	18	1.8
Dry Grocery	522	52.2
Fresh Produce	60	6
Frozen Foods	85	8.5
General Merchandise	9	0.9
Health and Beauty	4	0.4
Non-Food Grocery	43	4.3
Packaged Meat	40	4

Notes:

store/week combination, we have quantity and price data for at least 250 of our 1000 sample products. ii) For each good/week combination, we have quantity and price data for at least for at least 50 stores. These selection criteria ensure that we only focus on store/goods/weeks where we have sufficient data to reliably estimate the good-time means (10), and store-time means (12) in the first and second stages.

3.3 Results

We first illustrate our approach with two example states: Minnesota and Arizona, while the following section provides results for the nationwide sample.

Empirical Auto-covariance Function Figure 1 shows the auto-covariance function for the store-good component out to 100 lags.

Both states show a similar-shaped auto covariance function. There is a sharp decrease at the first lag, suggesting the presence of a large temporary component in prices. There is a further decrease at the second lag (which is more pronounce for Arizona than Minnesota) which suggests the presence of an MA(1) component. The function then decreases gently, suggesting the presence of a very persistent AR(1) component. The function does not decline all the way to zero, even after nearly two years, suggesting the presence of fixed effects in prices. Based on the shape of this function we use an AR(1) + MA(1) + FE as our baseline

Table 2: Data by department

Yogurt	107	10.7
Carbonated Beverages	93	9.3
Fresh Produce	60	6
Bread and Baked Goods	53	5.3
Pizza / Snacks / Hors doerves - Frozen	44	4.4
Milk	36	3.6
Vegetables - Canned	34	3.4
Soft Drinks - Non Carbonated	33	3.3
Soup	33	3.3
Candy	32	3.2
Cereal	30	3
Fresh Meat	30	3
Snacks	30	3
Cheese	29	2.9
Paper Products	28	2.8
Breakfast Food	23	2.3
Crackers	21	2.1
Dressings / Salads / Prep Foods - Deli	18	1.8
Prepared Food - Dry Mixes	18	1.8
Pasta	17	1.7

Notes:

model for the store-good component.

Figure 2 shows the auto-covariance function for the store component out to 100 lags. The function suggests that almost all of the variation in the store component is very persistent in nature and that the $AR(1) + MA(1) + FE$ fits the data well.

Baseline Parameter Estimates Table 3 shows the parameter estimates from the baseline model. The estimates show that almost all of the store component is persistent. Separately identifying the $AR(1)$ and the FE is difficult since the ACF is so flat. For AZ it is all loaded into the $AR(1)$ component which has a very high AR parameter. For MN more is loaded directly onto the FE. In both cases, less than 5% of the variation is transitory. For the store-good component about 2/3 of the variation is transitory. The remaining 1/3 is persistent.

This model fits the data very well. Figure X in Appendix A shows the observed and estimated

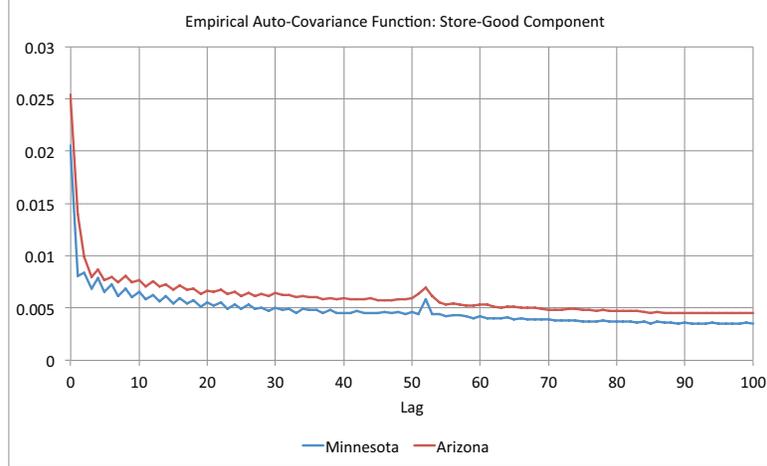


Figure 1: Impact of friction on steady-state

Notes:

auto covariance functions for the store-good component. The store component fit is so good that the model and data lines are almost indistinguishable.

Table 3: Data by department

	Store component		Store-good component	
	MN	AZ	MN	AZ
<u>Parameter estimates</u>				
θ : MA(1) parameter	0.152	0.209	0.049	0.377
σ_ε^2 : Variance of transitory shock	1.45E-04	1.47E-04	1.30E-02	1.48E-02
ρ : AR(1) parameter	0.983	0.999	0.966	0.977
σ_η^2 : Variance of persistent shock	2.32E-05	7.93E-06	2.71E-04	1.99E-04
σ_α^2 : Variance of fixed effect	2.57E-03	1.35E-10	3.54E-03	4.11E-03
<u>Variance decomposition</u>				
MA(1)	4%	3%	63%	67%
AR(1)	20%	97%	20%	17%
FE	76%	0%	17%	16%
Fraction due to AR(1) + FE	96%	97%	37%	33%

Notes:

Nationwide Results We run the same procedure state by state, for broader coverage as well as to see whether there is significant variation across states. Figure 3 shows the range of cross-sectional dispersion in the store and store-good components, respectively, confirming our earlier numbers. The standard deviation of the store component averages to 6% across

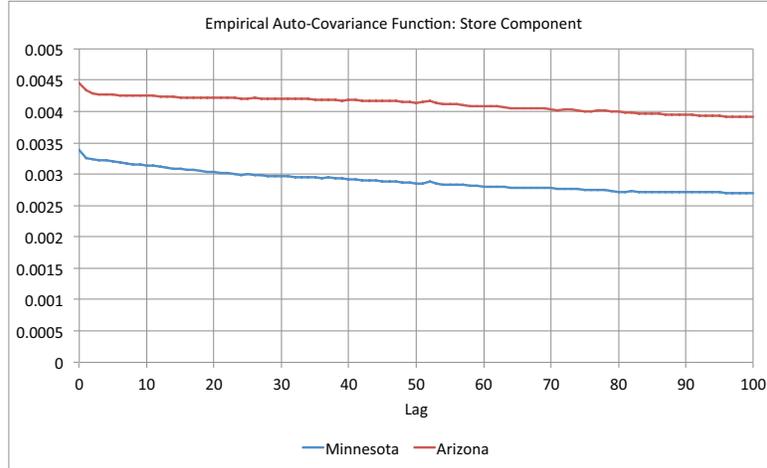


Figure 2: Impact of friction on steady-state

Notes:

states, while the standard deviation of the store-good component to 14%.

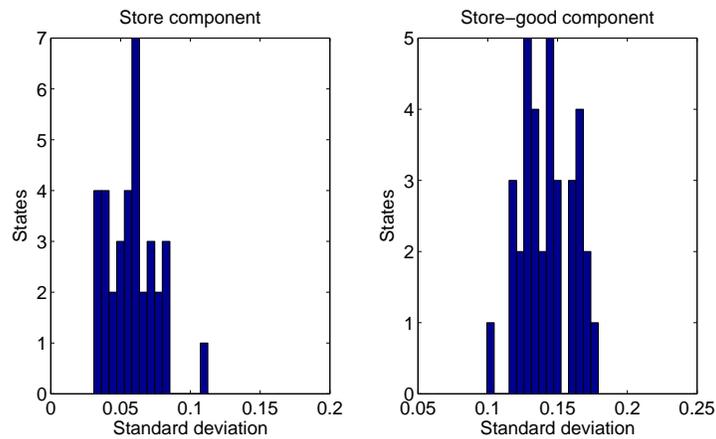


Figure 3: Histogram of standard deviations by state

Notes:

The variance decomposition shows that the share of persistent variation in the store component averages to 94% across states, while the share of persistent variation in the store-good component to 37%.

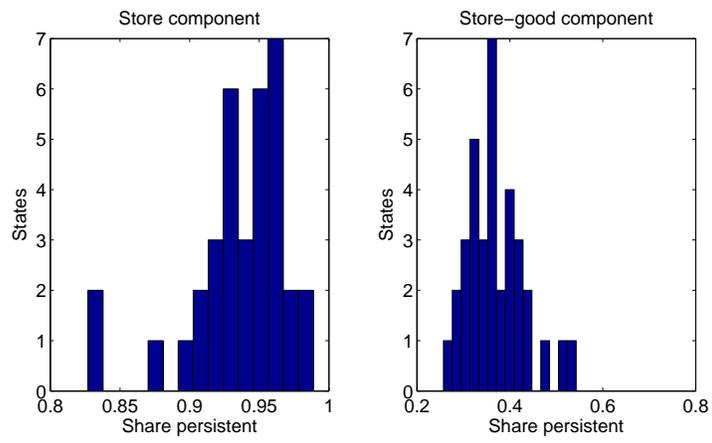


Figure 4: Histogram of share of persistent effect by state

Notes:

4 Theory

This section proposes a multi-product search model, which formalizes the idea of heterogeneity in households' willingness to shop around for a basket of goods. We build on Burdett and Judd (1983).

Consider a market for two goods, with a continuum of stores offering both for sale. There are measure m households per store, and each household can purchase either both goods, only one of them, or nothing. The households come in two types: the busy and the frugal, with fraction μ_b busy and $1 - \mu_b$ frugal. The households differ in their value of time, captured as follows. First, households do not have perfect information on prices in the market, but learn about them probabilistically: each household learns the prices of a single store with probability λ_i and the prices of two stores with probability $1 - \lambda_i$, where $i = \{b, f\}$. The busy are more likely to learn only one set of prices than the frugal, so $\lambda_b > \lambda_f$.⁴ Second, the busy can only visit one store to make their purchases, while the frugal can visit multiple stores (among those they have price information on). And finally, busy households have a higher willingness to pay for the goods than the frugal, $u_b > u_f$.

The household problem Each household learns either one or two sets of prices (p_1, p_2) . The frugal households maximize the objective

$$(u_f - p_1)\mathbf{1}_{\{\text{buy good 1 at } p_1\}} + (u_f - p_2)\mathbf{1}_{\{\text{buy good 2 at } p_2\}}. \quad (14)$$

If the household learns only one set of prices (p_1, p_2) , then these are prices available to the household to purchase goods at. Comparing the prices to their willingness to pay u_f , the household chooses whether to purchase both, only one, or neither good from the store in question. If the household learns two sets of prices, they have more options. In particular, because the frugal households can visit multiple stores, they can purchase each good where it is cheapest. They determine the lowest price for each good separately, and purchase unless the price exceeds their willingness to pay, visiting both stores if necessary. Due to

⁴One could think of busy households as observing fewer advertisements, for example, or having less acquired information based on having shopped in fewer stores in the past.

their ability to multi-stop shop, the frugal thus tend to make decisions independently across goods.

The busy households maximize the objective

$$(u_b - p_1)1_{\{\text{buy good 1 at } p_1\}} + (u_b - p_2)1_{\{\text{buy good 2 at } p_2\}}. \quad (15)$$

If the busy household only learns one set of prices, their situation is identical to the frugal. If the busy learn two sets of prices, their options differ from the frugal, however, because they can only visit one store to make purchases. Because they can only visit one store, busy households tend to make their purchase decisions minimizing the basket price $p_1 + p_2$ across stores, rather than minimizing individual prices. They will not pay more than u_f for either good, however, as they can also choose to purchase only one, or neither, of the two goods. But overall, as one-stop shoppers, the busy tend to make decisions based on basket prices. Note that it follows that households visiting multiple stores generally pay less per good than households only visiting one.

The store problem Each store sets its prices for the two goods to maximize profit $\pi(p_1, p_2)$, taking as given the distributions of prices in the market $F(p_1, p_2)$. We consider symmetric equilibria, denoting the marginal distributions for goods $i = 1, 2$ as $F_p(p_i)$. We also denote the distribution of basket prices $q = p_1 + p_2$ as $F_q(q)$. We return to spell out the store problems in the next section.

Definition of equilibrium An equilibrium is a level of profit π^* and distribution of prices $F(p_1, p_2)$ such that: i) Given the one or two sets of prices they observe, frugal households visit either one or two sets of stores to maximize (14). ii) Given the one or two sets of prices they observe, busy households visit one store to maximize (15). iii) $\pi(p_1, p_2) = \pi^*$ on the support of F and $\pi(p_1, p_2) \leq \pi^*$ everywhere.

With the model setup laid out, we next proceed to constructing equilibria. We begin with an equilibrium which demonstrates how multi-good search differs from one-good search, even if households continue one-stop shop, i.e., make all their purchases from a single store. In

this equilibrium only the busy households are active, and as a consequence both households and stores focus mainly on the basket price, bringing the problem closer to the familiar one-dimensional case. This case also provides a useful starting point for turning to multi-stop shopping in Section 4.2.

4.1 Equilibrium 1: one-stop shopping

Consider an equilibrium where stores target the busy households only – perhaps because the willingness to pay of the frugal is low, or because there are few frugal households in the market. Because the busy tend to focus on the prices of baskets rather than individual goods, the problem becomes similar to one-dimensional search, with some caveats, providing an introduction into two dimensions.

Given the focus of the busy is on basket prices, consider a distribution of basket prices q on an interval $[\underline{q}, \bar{q}]$, along the lines of Burdett and Judd (1983).

The top of the distribution \bar{q} cannot exceed $2u_b$, because this is the most a busy household is willing to pay for the basket. A store setting a higher basket price would have zero sales revenue, and profits, while equilibrium profits are positive. In fact, the top of the distribution must equal $2u_b$: If the top were to lie below this level, then a store offering this highest price would have incentive to deviate to $2u_b$, because doing so would increase revenue due to the higher price, without reducing demand (altering any buyer’s decision to buy). With the multi-product nature of the problem in mind, note that the prices of the individual goods cannot exceed u_b either, as the households could also choose to purchase only part of the basket. It would not be profit-maximizing for a store to forgo sales in this way.

Store profits at the top of the basket distribution equal $m\mu_b\lambda_b2u_b$, where there are m households per store, with fraction μ_b of them busy. A store pricing at the top of the price distribution only sells to those households learning a single price – fraction λ_b of households – and makes a profit $2u_b$ per customer.⁵ In equilibrium stores must make equal profits over

⁵We assume the store’s costs of supplying the goods are zero, or sunk.

the distribution of basket prices, and in particular these profits must equal those at the top:

$$\pi^* = m\mu_b\lambda_b 2u_b. \quad (16)$$

Note that equilibrium profits are positive, $\pi^* > 0$.

Profits for basket prices just below the top, for $q \leq 2u_b$, read

$$\pi(q) = m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q))]q. \quad (17)$$

The store continues to sell to all busy households learning a single set of prices only, but now also sells to households learning two sets of prices. For the latter: the store's prices are observed by $m\mu_b(1 - \lambda_b)2$ such households, each of whom purchase from this store only if their other price exceeds q , which happens with probability $1 - F_q(q)$. These profits must be constant on the support of the basket price distribution, implying that $\pi(q) \equiv \pi^*$ on the support. It follows that the distribution $F_q(q)$ must take form

$$F_q(q) = 1 + \frac{\lambda_b}{2(1 - \lambda_b)}\left(1 - \frac{2u_b}{q}\right), \quad (18)$$

with lower bound $\underline{q} = \frac{\lambda_b}{\lambda_b + 2(1 - \lambda_b)}2u_b$. Offering a basket price below the bottom of the distribution would not be profit-maximizing for stores, because such stores could raise their prices to \underline{q} without reducing demand.⁶

Finally, in order for the frugal households to remain out of the market, the price of each of the individual goods must exceed their willingness to pay, u_f . This implies a lower bound on the basket prices as well, $u_f + u_b$, as once the price of one good falls below u_f , the frugal enter the market to buy that good even if the other good were priced at the highest possible price, u_b .

⁶Here we have assumed a continuous support for the basket price distribution. If the distribution had a masspoint, it would be profitable for an individual store to deviate slightly below the masspoint: The negative effect of the price decline on profits would be more than offset by the increase in quantity sold. This means that the support must be continuous.

Equilibrium outcomes To illustrate equilibrium outcomes, Figure 5 plots the support of the price distribution, with the prices p_1, p_2 of the individual goods on the axis. Note that the downward sloping lines in the figure correspond to different price schedules (p_1, p_2) with the same basket price $p_1 + p_2$, with the line intersecting the axis at the basket price. The 45° line, on the other hand, corresponds to baskets with different basket prices q but equal prices for both goods $q/2$. Equilibrium basket prices range from \underline{q} to $2u_b$, but with individual prices constrained from above by u_b . The lowest basket price, \underline{q} , also must exceed $u_b + u_f$, to keep the frugal households out. The range of the support is denoted in the figure in red. Note, however, that while the equilibrium pins down the distribution of basket prices, it does not pin down the distributions of individual goods' prices. Thus the support need not cover the entire triangle in the graph, but could, for example, lie solely on the 45° line (or some symmetric region around it).

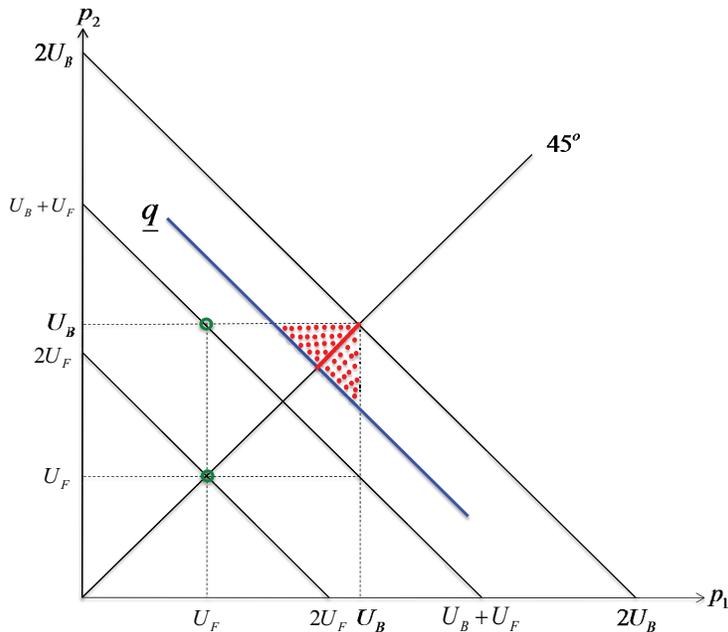


Figure 5: One-stop shopping

Notes: The figure plots the support of the price distribution, with the prices p_1, p_2 of the individual goods on the axis. The basket prices range from $2u_b$ down to $\underline{q} (> u_b + u_f)$.

Here stores are filled with busy households only, and all purchases take the form of baskets. The busy do not shop around, but make all their purchases in the one store with the lowest

basket price they know of. As long as some of these households learn more prices than others, there is dispersion in basket prices across stores: households with more information buy their basket where it is cheaper, pushing down the bottom of the price distribution, while the rest buy where they can, keeping the top of the distribution high. In the extremes: if all busy only learn one set of prices, stores price to their willingness to pay everywhere, whereas the more busy that learn multiple prices, the lower prices become. Of course eventually, as basket prices fall enough, this equilibrium disappears as the frugal households enter.

Note that the distribution of basket prices in the model corresponds conceptually to the distribution of store price levels in the empirical exercise. The model is thus able to generate this feature of the data, as long as some households learn multiple sets of prices. The model also accommodates dispersion in relative prices: for a given basket price q , a range of individual goods prices are generally consistent with $p_1 + p_2 = q$.

Existence of equilibrium Would stores prefer to deviate from this strategy of targeting the busy households only? The two deviations to be concerned about, in particular, are: (i) cutting the price of one good enough to sell it to the frugal also, pricing the two goods at (u_f, u_b) , or (u_b, u_f) , and (ii) cutting the price of both goods enough to sell them to the frugal also, pricing at (u_f, u_f) . Both price cuts would lead to an increase in sales volume, in return for the lower price.⁷

To rule out the first deviation, we need that

$$\underbrace{\pi^*}_{\text{equilibrium profit}} \geq \underbrace{\mu_b(\lambda_b + 2(1 - \lambda_b))(u_b + u_f)}_{\text{profit from deviation to } (u_b, u_f) \text{ or } (u_f, u_b)} + \underbrace{(1 - \mu_b)(\lambda_f + 2(1 - \lambda_f))u_f}_{\text{frugal buy singles}}. \quad (19)$$

This deviation would make the store the only one pricing one good low enough for frugal households to purchase also – hence all frugal households who learn its prices would purchase *that good* at the price u_f . But it would also make the store the lowest-basket-price store in the market – and hence all busy households who learn its prices would purchase a basket at

⁷Why not other deviations? Note that when deviating, there is no reason to cut prices lower than this, because doing so has no effect on demand.

the price $u_b + u_f$.

To rule out the second deviation, we need that

$$\underbrace{\pi^*}_{\text{equilibrium profit}} \geq \underbrace{\mu_b(\lambda_b + 2(1 - \lambda_b))2u_f}_{\text{busy buy baskets}} + \underbrace{(1 - \mu_b)(\lambda_f + 2(1 - \lambda_f))2u_f}_{\text{frugal buy baskets}}. \quad (20)$$

profit from deviation to (u_f, u_f)

This deviation would make the store the only one pricing its goods low enough for frugal households to purchase – hence all frugal households who learn its prices would purchase *both goods* at the price u_f . But it would also make the store the lowest-basket-price store in the market – and hence all busy households who learn its prices would purchase a basket at the price $2u_f$.

It turns out that the first condition implies the second, as well as the earlier requirement that $\underline{q} \geq u_f + u_b$. Rewriting condition (19) as (21) below, and defining $\rho = u_f/u_b$ and $\theta = (1 - \mu_b)/\mu_b$, we arrive at the following result:

Proposition 1. (*Existence of Equilibrium 1*) *If the following condition holds:*

$$\rho \leq \frac{3\lambda_b - 2}{2 - \lambda_b + \theta(2 - \lambda_f)}, \quad (21)$$

there exists a unique $\underline{q} \in [u_b + u_f, 2u_b]$, with the basket distributions given above, where conditions (19) and (20) hold.

This means that: First, the equilibrium exists only if busy households are not too likely to learn multiple prices (if $\lambda_b < \frac{2}{3}$). Their learning multiple prices makes deviating more attractive, as competition drives down the bottom of the price distribution. Second, the equilibrium exists only if frugal households have a low enough willingness to pay relative to the busy. Their having a high willingness to pay also makes deviating more attractive, by making the frugal more profitable customers. Finally, and for the same reason, the equilibrium exists only if the share of frugal households in the market remains small enough.

Distribution of prices The distributions take a relatively simple form in the model, allowing solving for moments in closed form. We can thus construct model counterparts to

those measures of price dispersion considered in the empirical part.

The model counterpart to the distribution of store price levels is the distribution of basket prices. The density of this distribution, from the expression in (18), takes form

$$f_q(q) = \frac{\lambda_b}{1 - \lambda_b} \frac{u_b}{q^2}. \quad (22)$$

Using this, we can solve for the mean and variance, arriving at the following:

Proposition 2. *(Store component) In Equilibrium 1, the distribution of basket prices q has mean $\mu_q = \frac{\lambda_b u_b}{1 - \lambda_b} \log(\bar{q}/\underline{q})$, and variance $\sigma_q^2 = \frac{\lambda_b u_b}{1 - \lambda_b} (\bar{q} - 2\mu_q \log \bar{q} - \mu_q^2/\bar{q} - \underline{q} + 2\mu_q \log \underline{q} + \mu_q^2/\underline{q})$.*

To develop a corresponding measure of the dispersion in relative prices, we must come to terms with the fact that if households always purchase baskets, the equilibrium does not pin down the distribution of individual goods' prices. A store with a given basket price may choose to price all goods equally expensive, or balance a lower price one good with a higher one on another. If all stores behave the same, we will observe little relative price dispersion in the first case, and more in the second. Even though we cannot pin down the standard deviation exactly, we can derive bounds on it, however, using the fact that individual prices cannot exceed u_b (see Figure 5).

We know that, conditional on basket price q , the minimum standard deviation of individual prices is zero, which happens if all stores price both goods the same. Conversely, conditional on basket price q , the maximum standard deviation of individual prices is $u_b - q/2$, which happens if all stores price at the extremes of $(u_b, q - u_b)$ or $(q - u_b, u_b)$. Integrating over basket prices, using the density (22), we arrive at the following bounds on relative price dispersion:

Proposition 3. *(Store-good component) In Equilibrium 1, the distribution of individual prices p around the basket price q has mean $\mu_{p|q} = 0$ and variance*

$$0 \leq \sigma_{p|q}^2 \leq \frac{\lambda_b u_b}{1 - \lambda_b} [\bar{q}/4 - u_b \log \bar{q} - u_b^2/\bar{q} - \underline{q}/4 + u_b \log \underline{q} + u_b^2/\underline{q}] \quad (23)$$

We can also derive the distribution of purchased (as opposed to posted) prices. For a given

$q \in [\underline{q}, 2U_B]$, we observe the measure of purchases

$$m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q))] \times f_q(q) = \frac{\pi^*}{q} \times \frac{\lambda_b}{1 - \lambda_b} \frac{u_b}{q^2} = \frac{\lambda_b}{1 - \lambda_b} \frac{u_b \pi^*}{q^3}, \quad (24)$$

where we have used the fact that profits equal those at the top of the distribution. Integrating over $[\underline{q}, 2u_b]$ yields total purchases $m\mu_b$. With this, we can write the density of purchased prices as

$$\hat{f}_q(q) = \frac{\lambda_b^2}{1 - \lambda_b} \frac{2u_b^2}{q^3}, \quad (25)$$

for $q \in [\underline{q}, 2u_b]$.

Proposition 4. *(The average purchased price is lower than the average posted price) The mean of purchased prices, $\mu_{\hat{q}} = \lambda_b 2u_b$, and the mean of listed prices, $\mu_q = \frac{\lambda_b}{1 - \lambda_b} \ln(\frac{2 - \lambda_b}{\lambda_b}) u_b$, satisfy $\mu_{\hat{q}} < \mu_q$, for $\lambda_b \in [0, 1)$, with the difference strictly decreasing in λ_b , and $\mu_{\hat{q}} = \mu_q$, for $\lambda_b = 1$.*

The average purchased basket price is lower than the average posted basket price, whenever $\lambda_b < 1$. If some households learn multiple sets of prices, there is dispersion in basket prices, and those households buy at the cheaper price. Thus the distribution of purchased prices has more weight in the bottom than the distribution of posted prices, and the mean of purchased prices is lower than that of posted prices.⁸

While instructive as a starting point for moving from one-product toward multi-product search, this equilibrium does not capture the heterogeneity in household shopping behavior that we originally set out to model, as the frugal households remain inactive. The equilibrium does feature dispersion in basket prices, as well as allowing dispersion in relative prices, but *does not necessarily imply* dispersion in relative prices. To generate this feature, we turn to an equilibrium where the frugal households also become active, in the next section.

⁸In principle we can also compare households getting more and less information: Those households getting only one offer generate purchase price distribution equivalent to the posted price distribution $f_q(q)$. Those households getting two offers generate $(\hat{f}_q(q) - \lambda_b f_q(q))/(1 - \lambda_b)$. Similarly, the means for these households equal μ_q and $(\mu_{\hat{q}} - \lambda_b \mu_q)/(1 - \lambda_b)$. Of course, in practise we seldom observe the information set of agents.

4.2 Equilibrium 2: multi-stop shopping

Consider an equilibrium where some stores continue to target the busy households only, while others recognize the opportunity to boost sales by tapping into the market of frugal households also. Doing so becomes attractive if the willingness to pay of the frugal is high, or if there are many frugal households in the market. The mix of busy and frugal households in the market now means that while some market participants focus on baskets, others focus on individual goods, and stores will need to consider both in setting their prices. To maintain basket prices at relatively high levels, the natural way for a store to implement this is to cut the price of one good enough to make it attractive to frugal households, while keeping the price of the other good high.

We now look for a basket price distribution on the union of two closed intervals $[q, u_b + u_f] \cup [q^*, 2u_b]$, where the high range corresponds to stores targeting busy households only, and the low ones targeting the frugal only.

The top of the basket distribution has to equal $2u_b$ for the same reasons as before, with both individual goods' prices at most u_b . Stores pricing at the top of the distribution make the same profits as above $m\mu_b\lambda_b 2u_b$, and thus equilibrium profits equal $\pi^* = m\mu_b\lambda_b 2u_b$. And as before, profits immediately below the top, for $q \leq 2u_b$, read $\pi(q) = m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q))]q$. For profits to remain constant over this range, the cumulative distribution again must take the form (18).

The reason to expect a gap in the support of the distribution is that as stores reduce their basket prices from $2u_b$, they eventually reach the point where the frugal start to buy, where the customer base of the store changes in a discontinuous way: from only the busy buying to both the busy and frugal buying. If prices change continuously as this happens, profits will change discontinuously as a result – not consistent with equal profits over the range of prices offered.

The highest basket price at which the frugal start buying is $u_b + u_f$. Here the store prices one of the two goods just low enough to get the frugal to buy, at u_f , and while pricing the

distribution (18). Moreover, because the stores pricing below this gap will price either of the two goods below u_f – with half of stores choosing good one and half good two – we also have $2F_p(u_f) = F_q(q^*)$. Using these two observations to write an equation for q^* , we arrive at

$$\pi^* = m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q^*))](u_f + u_b) + m(1 - \mu_b)[\lambda_f + 2(1 - \lambda_f)(1 - F_q(q^*))]u_f. \quad (27)$$

More generally, stores selling baskets to busy households at price q and single goods to frugal households at price p , have profits

$$\pi(q, p) \equiv m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q))]q + m(1 - \mu_b)[\lambda_f + 2(1 - \lambda_f)(1 - F_p(p))]p, \quad (28)$$

where $q \in [\underline{q}, u_b + u_f]$ and $p \in [\underline{p}, u_f]$. For such a store to be indifferent over a range of prices (q, p) , these profits need to remain constant over this range. As stores can generally vary the two prices independently, it follows that profits must remain constant over the range of prices offered *on both baskets and single goods separately*. In other words, profits on busy households buying baskets $\pi^b(q) \equiv m\mu_b[\lambda_b + 2(1 - \lambda_b)(1 - F_q(q))]q$ satisfy $\pi^b(q) \equiv \pi^b(u_b + u_f)$ for basket prices $q \in [\underline{q}, u_b + u_f]$, and profits on frugal households buying singles $\pi^f(p) \equiv m(1 - \mu_b)[\lambda_f + 2(1 - \lambda_f)(1 - F_p(p))]p$ satisfy $\pi^f(p) \equiv \pi^f(u_f)$ for prices $p \in [\underline{p}, u_f]$.

This implies the price distributions

$$F_q(q) = 1 + \frac{1}{2(1 - \lambda_b)} \left(1 - \frac{\pi^b(u_b + u_f)}{m\mu_b} \frac{1}{q}\right), \text{ for } q \in [\underline{q}, u_b + u_f], \quad (29)$$

$$F_p(p) = 1 + \frac{1}{2(1 - \lambda_f)} \left(1 - \frac{\pi^f(u_f)}{m(1 - \mu_b)} \frac{1}{p}\right), \text{ for } p \in [\underline{p}, u_f]. \quad (30)$$

Equilibrium outcomes Figure 6 illustrates equilibrium outcomes by plotting the support of the distribution, with the prices of the individual goods on the axis. The figure is similar to Figure 5, except in this case equilibrium basket prices range from \underline{q} to $u_b + u_f$ and then from q^* to $2u_b$. The lowest basket price \underline{q} must exceed $2u_f$, to keep the frugal from purchasing baskets from the same store. In this case the equilibrium again pins down the distribution of

basket prices, but now it also pins down the distribution of single goods' prices in the range where the frugal purchase single goods, $p \in [p, u_f]$.

Here stores come in two types. Some are filled with busy households, with all purchases taking the form of baskets. Others have a mixed clientele, where along with the busy households buying baskets, there are frugal households buying single goods. The latter stores tend to be less expensive than the former. Households are divided into busy one-stop shoppers, who pick the cheaper store among the options they have and buy everything there, and the frugal multi-stop shoppers, who purchase their goods from multiple stores. The multi-stop shoppers end up paying less for their basket, as the maximum basket price for frugal households is $2u_f$, while the busy always pay more for the basket. Note that here the frugal only purchase a basket if they learn the prices of two stores with different goods as their cheap good, otherwise they purchase a partial basket only.

Note that the model now implies a strictly positive degree of relative price dispersion. In the range where multi-stop shopping takes places, half the stores price good 1 cheap and good 2 expensive, and half the opposite. No store prices the goods equally expensive, however.

Existence of equilibrium Would stores prefer to deviate from these pricing policies? The deviation to be concerned about is cutting the prices of both goods enough to sell them to the frugal, pricing both goods at u_f . To rule out this deviation, we need that

$$\underbrace{\pi^*}_{\text{equilibrium profit}} \geq \underbrace{\mu_b(\lambda_b + 2(1 - \lambda_b))2u_f}_{\text{busy buy baskets}} + \underbrace{(1 - \mu_b)(\lambda_f + 2(1 - \lambda_f)(1 - F_p(u_f)))2u_f}_{\text{frugal buy singles or baskets}}. \quad (31)$$

profit from deviation to (u_f, u_f)

This deviation would make the store the only one pricing its both goods low enough for frugal households to purchase, but at a higher price than the cheaper goods are generally offered in equilibrium. As a result, frugal households would purchase each good either if the household only learned one set of prices, or if the other store priced the good above u_f . The deviation would make the store the lowest-basket-price store in the market, and hence all busy households learning its prices would purchase a basket at the price $2u_f$.

We again arrive at a proposition for existence of equilibrium, which places constraints the

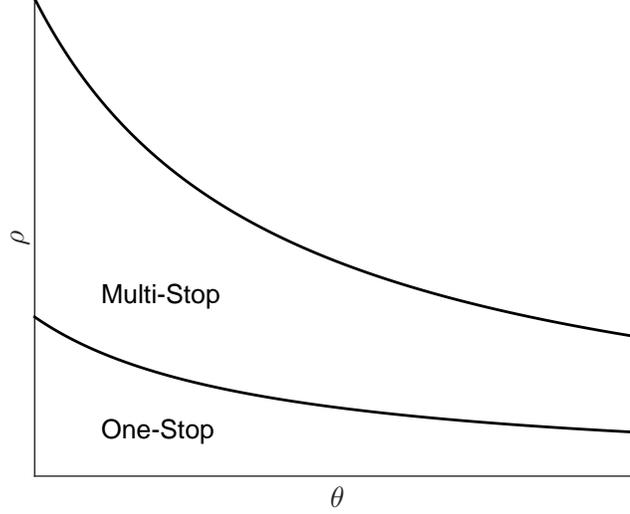


Figure 7: Equilibrium type

Notes: The figure illustrates the range of the parameter space where Equilibrium 1: One-stop shopping, and Equilibrium 2: Multi-stop shopping occur.

parameter values that come to question.

Proposition 5. (*Existence of Equilibrium 2*) *If the following conditions hold:*

$$\rho \geq \frac{3\lambda_b - 2}{2 - \lambda_b + \theta(2 - \lambda_f)}, \quad (32)$$

$$\rho \leq \frac{\lambda_b}{2 - \lambda_b + \theta \left[1 + \frac{1 - \lambda_f}{2(1 - \lambda_b)} \frac{\lambda_b(1 - \rho) - \theta\rho}{1 + \rho + \theta\rho \frac{1 - \lambda_f}{2(1 - \lambda_b)}} \right]}, \quad (33)$$

then there exists a unique $q^ \in [u_b + u_f, 2u_b]$ and $\underline{q} \in [2u_f, u_b + u_f]$, with the basket distributions given above, where condition (31) holds.*

To illustrate how the two relate to each other, Figure 7 plots the range of parameter values where the two equilibria occur in terms of θ and ρ . The upper bound in condition (33) can be shown to always lie above the lower bound in condition (32), with both curves decreasing. Comparing Propositions 5 and 1 shows that the two equilibria never occur for the same parameter values, as condition (21) is the reverse of condition (32). As the figure shows, the multi-stop shopping equilibrium occurs when willingness to pay, and share of market participants, of the frugal households is greater, making targeting them profitable for stores.

Proposition 6. (*Stores matter*) *The distribution of prices is not consistent with stores pricing the two goods independently.*

Note that the multi-product nature of the search problem matters for outcomes here, in that the equilibrium is not consistent with stores pricing the two goods independently. In the region of store prices where multi-stop shopping occurs, the stores cut one price but not the other, although they can choose which one to cut.

Distribution of prices The shape of the basket price distribution again allows solving for moments in closed form, although the separate regions make expressions longer. The density of the basket price distribution, from (29), takes form

$$f_q(q) = \begin{cases} \frac{\lambda_b u_b}{1-\lambda_b} \frac{1}{q^2} & \text{if } q^* \leq q \leq \bar{q}, \\ \frac{\lambda_b + 2(1-\lambda_b)(1-F_q(q^*))}{2(1-\lambda_b)} \frac{u_f + u_b}{q^2} & \text{if } \underline{q} \leq q \leq u_b + u_f. \end{cases} \quad (34)$$

Using this, we can solve for the mean and variance, arriving at the following:

Proposition 7. (*Store component*) *In Equilibrium 2, the distribution of q has mean and variance*

$$\mu_q = \frac{\lambda_b u_b}{1-\lambda_b} \log(\bar{q}/q^*) + \frac{\lambda_b + 2(1-\lambda_b)(1-F_q(q^*))}{2(1-\lambda_b)} (u_b + u_f) \log((u_b + u_f)/\underline{q}) \quad (35)$$

$$\begin{aligned} \sigma_q^2 &= \frac{\lambda_b u_b}{1-\lambda_b} [\bar{q} - 2\mu_q \log \bar{q} - \mu_q^2/\bar{q} - q^* + 2\mu_q \log(q^*) + \mu_q^2/q^*] \\ &+ \frac{\lambda_b + 2(1-\lambda_b)(1-F_q(q^*))}{2(1-\lambda_b)} (u_b + u_f) [u_b + u_f - 2\mu_q \log(u_b + u_f) - \mu_q^2/(u_b + u_f) \\ &- \underline{q} + 2\mu_q \log(\underline{q}) + \mu_q^2/\underline{q}]. \end{aligned} \quad (36)$$

Note that this multi-stop shopping equilibrium always features dispersion in relative prices (see Figure 6). To develop a measure of the dispersion in relative prices, note that the minimum standard deviation of individual prices, conditional on basket price q , can be

written:

$$\underline{\sigma}_{p|q} = \begin{cases} 0, & \text{for } q^* \leq q \leq 2u_b, \\ q/2 - u_f, & \text{for } \underline{q} \leq q \leq u_b + u_f. \end{cases} \quad (37)$$

Similarly, the maximum standard deviation of individual prices, conditional on basket price q , can be written:

$$\bar{\sigma}_{p|q} = \begin{cases} u_b - q/2, & \text{for } q^* \leq q \leq 2u_b, u_b + \underline{p} \leq q \leq u_b + u_f, \\ q/2 - \underline{p} & \text{for } \underline{q} \leq q \leq u_b + \underline{p}, \end{cases} \quad (38)$$

where the point $u_b + \underline{p}$ represents the corner in Figure 6.

Integrating over basket prices, using the density (34), we arrive at the following bounds on relative price dispersion:

Proposition 8. *In Equilibrium 2, the distribution of individual prices p around the basket price q has mean $\mu_{p|q} = 0$ and variance*

$$\sigma_{p|q}^2 \geq \frac{\lambda_b + 2(1 - \lambda_b)(1 - F_q(q^*))}{2(1 - \lambda_b)} (u_b + u_f) \Big|_{\underline{q}}^{u_b + u_f} [q/4 - u_f \log(q) - u_f^2/q] \quad (39)$$

$$\begin{aligned} \sigma_{p|q}^2 \leq & \frac{\lambda_b u_b}{1 - \lambda_b} \Big|_{q^*}^{\bar{q}} [q/4 - u_b \log(q) - u_b^2/q] \\ & + \frac{\lambda_b + 2(1 - \lambda_b)(1 - F_q(q^*))}{2(1 - \lambda_b)} (u_b + u_f) \Big|_{u_b + \underline{p}}^{u_b + u_f} [q/4 - u_b \log(q) - u_b^2/q] \\ & + \frac{\lambda_b + 2(1 - \lambda_b)(1 - F_q(q^*))}{2(1 - \lambda_b)} (u_b + u_f) \Big|_{\underline{q}}^{u_b + \underline{p}} [q/4 - \underline{p} \log(q) - \underline{p}^2/q] \end{aligned} \quad (40)$$

5 Conclusions

TBW

References

- BLUNDELL, R., AND I. PRESTON (1998): “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, pp. 603–640.
- BURDETT, K., AND K. L. JUDD (1983): “Equilibrium Price Dispersion,” *Econometrica*, 51(4), 955–969.
- BURDETT, K., AND D. A. MALUEG (1981): “The Theory of Search for Several Goods,” *Journal of Economic Theory*, 24, 362–376.
- CARLSON, J. A., AND R. P. MCAFEE (1984): “Joint Search for Several Goods,” *Journal of Economic Theory*, 32, 337–345.
- EDEN, B. (2013): “Price dispersion and demand uncertainty: Evidence from us scanner data,” Mimeo, Vanderbilt University.
- GATTI, J. R. J. (1999): “Multi-Commodity Consumer Search,” *Journal of Economic Theory*, 86, 219–244.
- GOTTSCHALK, P., AND R. MOFFITT (1994): “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers*, pp. X–X.
- KAPLAN, G., AND G. MENZIO (2014): “The Morphology of Price Dispersion,” *International Economic Review*, pp. X–X.
- MCAFEE, R. P. (1995): “Multiproduct Equilibrium Price Dispersion,” *Journal of Economic Theory*, 67, 83–105.
- PRATT, J. W., D. A. WISE, AND R. ZECKHAUSER (1979): “Price Differences in Almost Competitive Markets,” *Quarterly Journal of Economics*, 93(2), 189–211.
- REINGANUM, J. F. (1979): “A Simple Model of Equilibrium Price Dispersion,” *Journal of Political Economy*, 87(4), 851–858.
- RHODES, A. (2015): “Multiproduct Retailing,” *Review of Economic Studies*, 82(1), 333–359.

- SALOP, S., AND J. STIGLITZ (1976): “Bargains and Ripoffs: A Model of Monopolistically Competitive Prices,” *Review of Economic Studies*, 44, 493–510.
- STIGLER, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69(3), 213–225.
- WILDE, L. L., AND A. SCHWARTZ (1979): “Equilibrium Comparison Shopping,” *Review of Economic Studies*, 46, 543–554.
- ZHOU, J. (2014): “Multiproduct Search and the Joint Search Effect,” *American Economic Review*, 104(9), 2918–39.