

THE MORPHOLOGY OF PRICE DISPERSION*

BY GREG KAPLAN AND GUIDO MENZIO¹

Princeton University and NBER, U.S.A.; University of Pennsylvania and NBER, U.S.A.

This article is a study of the shape and structure of the distribution of prices at which an identical good is sold in a given market and time period. We find that the typical price distribution is symmetric and leptokurtic, with a standard deviation between 19% and 36%. Only 10% of the variance of prices is due to variation in the expensiveness of the stores at which a good is sold, whereas the remaining 90% is due, in approximately equal parts, to differences in the average price of a good across equally expensive stores and to differences in the price of a good across transactions at the same store. We show that the distribution of prices that households pay for the same bundle of goods is approximately Normal, with a standard deviation between 9% and 14%. Half of this dispersion is due to differences in the expensiveness of the stores where households shop, whereas the other half is mostly due to differences in households' choices of which goods to purchase at which stores. We find that households with fewer employed members pay lower prices and do so by visiting a larger number of stores instead of by shopping more frequently.

mor.phol.o.gy (*noun*)

The branch of biology that deals with the form and structure of organisms without consideration of function.

1. INTRODUCTION

This article is a systematic study of the morphology of price dispersion, i.e., the shape and structure of the distribution of prices at which an identical good is sold in a given geographic market during a given period of time. Our goal is to provide macroeconomists with a set of stylized facts that can be used to develop and test theories of price dispersion and to calibrate models that emphasize deviations from the Law of One Price. We use data from the Kilts-Nielsen Consumer Panel (KNCP) data set, which contains price and quantity information for more than 300 million transactions by 50,000 households for over 1.4 million goods in 54 geographic markets during the period 2004–9.

Our approach views both data and theory through the lens of a decomposition of the cross-sectional variance of prices in a given market and time period. The decomposition expresses each transaction price as the sum of three sources: one that is specific to the store where the transaction took place, one that is specific to both the store where the transaction took place and the particular good that was traded, and one that is specific to the transaction itself. Such a decomposition can feasibly be implemented in the KNCP, since the data contain information on transactions for the same good at multiple stores, transactions for multiple goods at individual stores, and multiple transactions for the same good at individual stores. After investigating the shape of the typical price distribution, we use the decomposition to study three possible reasons for price dispersion. Does dispersion arise because some stores are more expensive than others

*Manuscript received February 2014; revised July 2014.

¹ This work uses data supplied by the Kilts-Nielsen Data Center at the University of Chicago Booth School of Business. Information on data access and availability is available at <http://research.chicagobooth.edu/nielsen>. We are grateful to the Griswold Center for Economic Policy Studies for providing financial support and to Mark Aguiar and two anonymous referees for helpful suggestions. Please address correspondence to: Greg Kaplan, Department of Economics, Fisher Hall, Princeton University Princeton, NJ 08544-1021. Phone: 609-258-4000. Fax: 609-258-6419. E-mail: gkaplan@princeton.edu.

on average, because the same good is sold at different prices at equally expensive stores, or because the same good is sold at the same store at different prices on different days?

We also investigate the shape and structure of dispersion in the average price paid by different households for the same bundle of goods. We do this by constructing household-specific price indexes, which is possible because the KNCP data link each transaction to a specific household. We measure the amount of dispersion in these price indexes and show that price dispersion does not wash out at the household level. We then use our decomposition to study the sources of price index dispersion across households, asking if some households pay less than others because they shop at cheaper stores, because they purchase cheaper goods at the same stores, or because they purchase the same goods at the same stores on different days.

We start to answer these questions by documenting the extent and shape of dispersion in the prices for an identical good in a given market and quarter. When we define goods by their Universal Product Code (UPC, i.e., barcode), the average standard deviation of normalized prices is 19%. Defining goods more broadly leads to more price dispersion. For our broadest definition, which aggregates into a single good all products that are identical except for their brand and size, the average standard deviation of normalized prices is 36%. We find that price dispersion is a prevalent phenomenon: While the amount of price dispersion does vary across goods, the standard deviation of prices is greater than 10% for more than 90% of goods, markets, and quarters. Moreover, for all definitions of a good, we find that the typical price distribution is symmetric, unimodal, and leptokurtic, meaning that the distribution has more mass around the mean and thicker tails than a Normal distribution with the same mean and variance. This contrasts with the typical features of the wage distribution, which is well known to be skewed with a long right tail.

To understand the sources of price dispersion, we decompose each transaction price into three components: (i) a store component, defined as the average price of all goods at the particular store where the transaction took place; (ii) a store-specific good component, defined as the average price of the particular good at that store relative to the average price of all goods at that store; and (iii) a transaction component, defined as the price of the good in that particular transaction relative to the average price of that good at that store. We find that the store component contributes approximately 10% to the overall variance of prices, the store-specific good component contributes between 25% and 45% depending on both how broadly we define a good and the selection criteria that we apply, and the transaction component accounts for the remaining variance. These findings are striking since they reveal that price dispersion does not occur primarily because some stores are cheap and some stores are expensive. Instead, price dispersion arises because, even among equally expensive stores, the average price of a specific good varies substantially, and even at a given store, the price of a specific good varies across transactions.

Our variance decomposition for transaction prices offers new perspectives on the relative importance of four dominant theories of price dispersion: (i) heterogeneity in the amenities offered by different stores; (ii) heterogeneity in the marginal costs faced by different stores; (iii) intertemporal price discrimination; and (iv) search frictions. Deeper understanding about each theory's relevance comes from comparing the predictions of each theory with the weight of each of the three components in our decomposition. Differences in amenities and differences in marginal costs are both unlikely to be important sources of price dispersion, since these differences would show up in the store component of prices, which accounts for only 10% of the overall variance of prices. In contrast, intertemporal price discrimination may be a quantitatively important source of price dispersion, since the time-variation in prices that stores use to discriminate among different types of buyers would be reflected in the variance of the transaction component of prices; this accounts for at least 45% of the overall variance of prices. Similarly, search frictions may be a quantitatively important source of price dispersion, since these frictions induce stores to randomize over prices, leading to variation in both the store and the store-specific good components of prices. Together, these two components account for at least 35% of the overall variance of prices.

We then ask whether some households are better than others at taking advantage of price dispersion and, if so, how they achieve this in their shopping behavior. For every household, we construct a price index by computing the ratio between the household's expenditures and the counterfactual expenditures had they purchased each good at its average price. We find that the distribution of price indexes in a given market and quarter is approximately Normal with a standard deviation between 9% and 14%, depending on how a good is defined. Hence, some of the dispersion in transaction prices disappears when goods are aggregated into bundles, but a substantial amount of variation remains. Indeed, a household at the 90th percentile of the price index distribution pays approximately 22% more than a household at the 10th percentile. Decomposing the variance of household price indexes into a store component, store-specific good component, and transaction component reveals that the store component accounts for approximately 50% of the variance of price indexes, whereas the store-specific good component accounts for 40%, and the transaction component accounts for only 10%.

Three revealing differences about household shopping processes emerge from a comparison of the variance decomposition for household price indexes with the decomposition for transaction prices: (i) there is less dispersion in household price indexes than in transaction prices; (ii) the store component accounts for a larger fraction of household price index dispersion than transaction price dispersion; and (iii) the transaction component accounts for a smaller fraction of household price index dispersion than transaction price dispersion. In Subsection 4.2, we describe a very simple shopping model that is consistent with these observations. The model relies on two key assumptions: (i) There is heterogeneity across households in the number of stores that they regularly visit, and (ii) all households are similar in their abilities to exploit temporary price reductions by visiting their chosen stores more frequently. We regress household price indexes on the number of shopping trips taken and the number of stores visited and find that visiting an additional store has a much larger effect on prices paid than visiting the same store more frequently, consistent with the simple shopping model that we describe.

We conclude the article by identifying some characteristics of households who pay lower prices. We find that household price indexes decline monotonically with age: The average price index for households older than 55 is between 3.5% and 4.5% lower than the average price index of households younger than 25. We also find that households with more nonemployed members pay less for the same bundle of goods than households whose members are all employed: The average price index for nonemployed households is between 1% and 4.5% lower than it is for employed households. Since older households and nonemployed households are likely to have a lower opportunity cost of time, these findings support the view that time may be a key input for the process of finding lower prices.

Our article contributes a systematic analysis of price dispersion to a literature that, until now, has mostly focused on case studies of particular goods. For example, Sorensen (2000) analyzed the features of the distribution of prices posted by different pharmacies for several drugs in two cities in upstate New York. Hong and Shum (2006) documented the distribution of prices posted by online booksellers for four academic textbooks. Moraga-González and Wildenbeest (2008) documented the distribution of prices posted by online sellers for several computer memory chips. Woodward and Hall (2012) documented price dispersion for mortgage brokerage services. Galenianos et al. (2012) documented and interpreted price dispersion for several illegal drugs. Baye et al. (2006) reviewed additional case studies of price dispersion. Although these case studies are methodologically important, they do not provide a general understanding of the shape and structure of price dispersion in the retail sector.

A recent and important exception to the lack of large-scale studies of price dispersion is Eden (2013), who uses transaction data from several supermarket chains in Chicago to relate the extent of price dispersion to characteristics of goods sold, such as the volatility of demand for the good, the expensiveness of the good, and the share of revenues associated with sales of the good. This study is a natural complement to ours, which focuses more on the 'typical' morphology of price dispersion.

Our article also contributes to a recent strand of literature documenting that different people pay different prices to achieve the same level of consumption. This literature was initiated by Aguiar and Hurst (2005), who observed that the difference in food consumption between retired and nonretired households is significantly smaller than the difference in food expenditures because retired households use home production to achieve more consumption per dollar of food expenditures. In follow-up work, Aguiar and Hurst (2007) used a subset of the KNCP data set to show that household price indexes decline with age. Our article extends those findings by showing that household price indexes are also lower for nonemployed households than for employed households. Moreover, our article documents the shape and structure of the overall dispersion in households' price indexes. We use these findings in a related paper (Kaplan and Menzio, 2014) to show that differences in shopping behavior between employed and unemployed buyers can be so strong as to generate multiple equilibria associated with different levels of employment and economic activity.

2. KILTS-NIELSEN CONSUMER PANEL DATA

2.1. *The Data.* The KNCP is a panel data set that tracks the shopping behavior of approximately 50,000 households over the period 2004–9. Households in the KNCP are drawn from 54 geographically dispersed markets, known as Scantrack markets, each of which roughly corresponds to a Metropolitan Statistical Area (MSA).² Demographic data on household members are collected at the time of entry into the panel and are then updated annually through a written survey during the fourth quarter of each year. The collected information includes age, education, marital status, employment, type of residence, and race.³

Households in the panel provide information about each of their shopping trips using a UPC, i.e., barcode, scanning device provided by Nielsen. When a panelist returns from a shopping trip, he uses the device to enter details about the trip, including the date and store where the purchases were made. The panelist then scans the barcode of each purchased good and enters the number of units purchased. The price of the good is recorded in one of two ways, depending on the store where the purchase took place. If the good was purchased at a store that Nielsen covers, the price is set automatically to the average price of the good at the store during the week when the purchase was made. If the good was purchased at a store that Nielsen does not cover, the price is directly entered by the panelist. Panelists are also asked to record whether the good was purchased using one of four types of deals: (i) store feature, (ii) store coupon, (iii) manufacturer coupon, or (iv) other deal. If the deal involved a coupon, the panelist is prompted to input its value.⁴

Nielsen offers households a variety of incentives to join and stay active in the panel. These include monthly prize drawings, gift points for weekly data transmission, and regular sweepstakes. Nielsen maintains a purchasing threshold, based on household size, that is used to filter out households who are unreliable reporters and, therefore, could bias results. The KNCP data contain only actively reporting households who meet this threshold over a 12-month period. This yields an annual panelist retention rate of approximately 80%.⁵ New households are

² Table A.1 in Appendix A.1 lists the 54 Scantrack markets, together with the fraction of observed expenditures that comes from each market. Six markets contain more than one MSA (Buffalo-Rochester, Hartford-New Haven, New Orleans-Mobile, Oklahoma City-Tulsa, Raleigh-Durham, and Salt Lake City-Boise) and one MSA is made up of three Scantrack markets (New York City). Our analysis uses data from all 54 markets from 2004 to 2009. However, for illustrative purposes and because of the sheer magnitude of the number of transactions in the data set, we will return throughout the article to one example market, Minneapolis, in one quarter, 2007:Q1.

³ Einav et al. (2010), Handbury (2012), and Bronnenberg et al. (2013) also use the KNCP. In their analysis of the shopping behavior of retired households, Aguiar and Hurst (2007) use a smaller version of the KNCP that covers only one market.

⁴ Starting in 2007 and continuing to 2008–9, the fraction of transactions recorded with a deal (predominantly a store feature) declined sharply from 30% to 24%. This change was due to a new method of capturing deals on newer devices given to incoming panelists. We have verified that our findings are not affected by this change.

⁵ Approximately, 20% of panelists are present for all six years, with a median duration of three years.

regularly added to the panel to replace attritions. Nielsen uses a stratified sampling design to ensure that the panel is demographically balanced. To make the sample representative of the U.S. population, Nielsen constructs projection weights; we apply these weights in our analyses.⁶

The products purchased by panelists are identified by UPC, of which the database contains around 1.4 million.⁷ These are grouped into 1,082 *product modules*, which are grouped into 113 *product groups*, which are then grouped into 10 *departments*. For each UPC, the data contain information on several product characteristics such as brand, size, packaging, and a rich set of product features. The recorded features differ by product module and depend on the nature of the product.⁸ The full set of departments, together with the number of product groups, product modules, and UPCs contained in each, is listed in Table A.2 in Appendix A. As an illustrative example, Table A.2 also reports the number of different UPCs purchased from each department in Minneapolis in 2007:Q1 and the fraction of overall expenditure coming from each department.⁹

The transactions in the KNCP data involve purchases from a variety of different types of retailers. About half of all expenditure comes from retailers in the Grocery category, and about a third comes from the Discount and Warehouse Club categories. The breakdown of expenditures by retailer is possible because the KNCP data contain retailer identifiers for every transaction. Moreover, for approximately one-third of all transactions, the data also contain identifiers for the store (within the retail chain) where the transaction took place. Since we are interested in assessing the fraction of price dispersion due to differences in the average prices at different stores, we only use transactions for which the store identifier is present.¹⁰

2.2. Defining a Good. We use the KNCP data to measure the extent of price dispersion, by which we mean the variation in the prices at which the same good is traded in a given market and time period, and to measure the extent of household price index dispersion, by which we mean the variation in the average price paid for the same bundle of goods by different households in a given market and time period. In order to carry out these measurements, we first need a definition for a good.

A natural starting point is to define a good by its UPC. However, there are reasons to group products with different UPCs together and define them as the same good. First, the same product may be given two or more UPCs for purely administrative reasons. Second, there may be different products that have different UPCs, but that are very close substitutes from the perspective of buyers. Consider the following examples: Bottles of ketchup that are produced

⁶ Sample weights are based on the following demographic characteristics: household size; household income; household head age; education, race, and occupation; and presence of children.

⁷ The KNCP data also include information on transactions for goods without barcodes (*magnet* transactions). We exclude these goods from our analysis since they are recorded for only a small nonrandom subset of households in the panel.

⁸ For example, the Catsup product module that contains Heinz Ketchup (a product we describe in detail in Subsection 3.1) has information on the following product features: size, type of container, type (regular vs. light vs. low fat, etc.), style (regular vs. fancy vs. hot & spicy vs. zesty garlic vs. onion, etc.), and organic status.

⁹ Nielsen estimates that approximately 30% of household consumption is accounted for by consumer panel data categories. Broda and Parker (2012) compared weekly spending in the KNCP and the Consumer Expenditure Survey (CEX). They found that the KNCP has about 30% of the expenditures of the CEX for a broad measure of nondurables and about 20% of the expenditures in the CEX across all consumption categories.

¹⁰ Table A.3 in Appendix A reports the fraction of expenditures coming from each type of retailer in 2007:Q1 and the fraction of expenditures coming from each type of retailer in 2007:Q1 for the subset of transactions with store identifiers. We cannot guarantee that the set of transactions for which we observe store identifiers is a random subset of all transactions. However, some of our calculations, such as measuring the overall amount of price dispersion in Subsection 3.2, do not require us to observe identifiers for individual stores. For these calculations, we have repeated the analysis using the full set of transactions and have found virtually identical results. Although the KNCP data do not technically distinguish prices entered by the panelist from those that come from point-of-sale data, we believe that the ‘covered’ stores from which Nielsen obtains point-of-sale price data are the same stores that have unique store identifiers. Since we restrict attention to these stores, we believe that the transactions included in our analysis mostly pertain to point-of-sale prices.

TABLE 1
DEFINITIONS OF A GOOD

Definition	Number of Products		Characteristics Held Fixed			
	in Database	Sold in Mpls, 2007:Q1	UPC	Brand	Size	Other Features
Department	10	10				
Product Group	113	113				
Product Module	1082	939				
Four definitions of a good:						
Brand & Size Aggregation	273,452	17,108				X
Brand Aggregation	606,969	25,201			X	X
Generic Brand Aggregation	817,483	30,446		X	X	X
UPC	1,424,453	36,104	X	X	X	X

NOTES: Number of products sold in Minneapolis in 2007:Q1 refers only to transactions at stores with a store identifier.

by the same manufacturer may be sold by different retailers under different generic brands and have different UPCs. Yet, these bottles of ketchup are likely to be perfect substitutes for buyers. Bottles of ketchup of different brands have different UPCs, but are likely to be very close substitutes for buyers. Bottles of ketchup of different sizes have different UPCs, but are also very likely to be close substitutes for buyers.

Based on the above observations and examples, we define four notions of a good, described below in order from narrowest to broadest.

UPC. A good is the set of products with the same UPC.

Generic Brand Aggregation. A good is the set of products that share the same features, the same size, and the same brand, but may have different UPCs. Since the KNCP assigns the same brand code to all generic brands (regardless of the retailer), this definition collects all generic brand products that are otherwise identical.

Brand Aggregation. A good is the set of products that share the same features and the same size, but may have different brands and different UPCs.

Brand and Size Aggregation. A good is the set of products that share the same features but may have different sizes, different brands, and different UPCs. Notice that for the first three definitions, prices are defined per unit of the good (e.g., the price of a 36-oz bottle of ketchup), whereas for this definition, prices must be defined per unit of measure (e.g., ounces or feet).

Table 1 reports the total number of goods in the KNCP database according to each definition as well as the actual number of goods for which we observe a transaction in Minneapolis in 2007:Q1.

When goods are defined by their UPC, there are 1,424,453 goods in the KNCP database, of which we observe 36,104 goods purchased in Minneapolis in 2007:Q1. The Generic Brand Aggregation reduces the number of goods in the KNCP database by 43% and the number purchased in Minneapolis in 2007:Q1 by 16%. The Brand Aggregation reduces the number of goods in the KNCP database by an additional 26% and the number purchased in Minneapolis 2007:Q1 by an additional 17%. Finally, the Brand and Size Aggregation reduces the number of goods in the KNCP database by an additional 55% and the number purchased in Minneapolis in 2007:Q1 by an additional 32%.

Since we wish to remain agnostic as to which is the most appropriate definition, we will measure price dispersion and price index dispersion for all four definitions. In cases where reporting results from more than one definition would unnecessarily hinder the presentation, we will use the Generic Brand Aggregation definition as a benchmark.¹¹

¹¹ In preliminary work, we have also considered other definitions of good (for example, aggregating sizes but not brands) and obtained similar findings.

2.3. *Data Organization and Selection.* For a given definition of a good, we organize the data by goods (j) and shopping trips (k). Each shopping trip is characterized by the household that went shopping $i(k)$, the market in which the household resides $m(i(k))$, the quarter in which the shopping trip took place $t(k)$, the store that was visited $s(k)$, and the retail chain to which the store belongs $r(k)$. We define $q_{j,k}$ as the quantity of good j purchased on shopping trip k and $P_{j,k}$ as the unit price paid for the good.

In order to compare price distributions for different goods and to compare price distributions for the same good in different markets and quarters, we need to normalize prices. We define the normalized price $p_{j,k}$ for good j on shopping trip k as the ratio between the price of the good in that shopping trip, $P_{j,k}$, and the average price of the good in all shopping trips that took place in the same market and quarter. Formally, $p_{j,k}$ is given by

$$(1) \quad p_{j,k} \equiv \frac{P_{j,k}}{\bar{P}_{j,m(k),t(k)}},$$

where $\bar{P}_{j,m,t}$ is given by

$$(2) \quad \bar{P}_{j,m,t} \equiv \frac{\sum_{j,k \in m \cap t} P_{j,k} q_{j,k}}{\sum_{j,k \in m \cap t} q_{j,k}},$$

and the notation $k \in m \cap t$ refers to $\{k : t(k) = t, m(i(k)) = m\}$, i.e., the set of shopping trips that took place in market m in quarter t .

We impose some selection criteria on transactions and goods. First, to ensure that we have a reliable estimate of the price distribution for each good, we only include a good j in a market m and quarter t if we observe a minimum number of transactions, n . For our baseline analysis, we set $n = 25$, but we have also repeated the analysis for $n \in \{15, 50\}$, and obtained similar results.

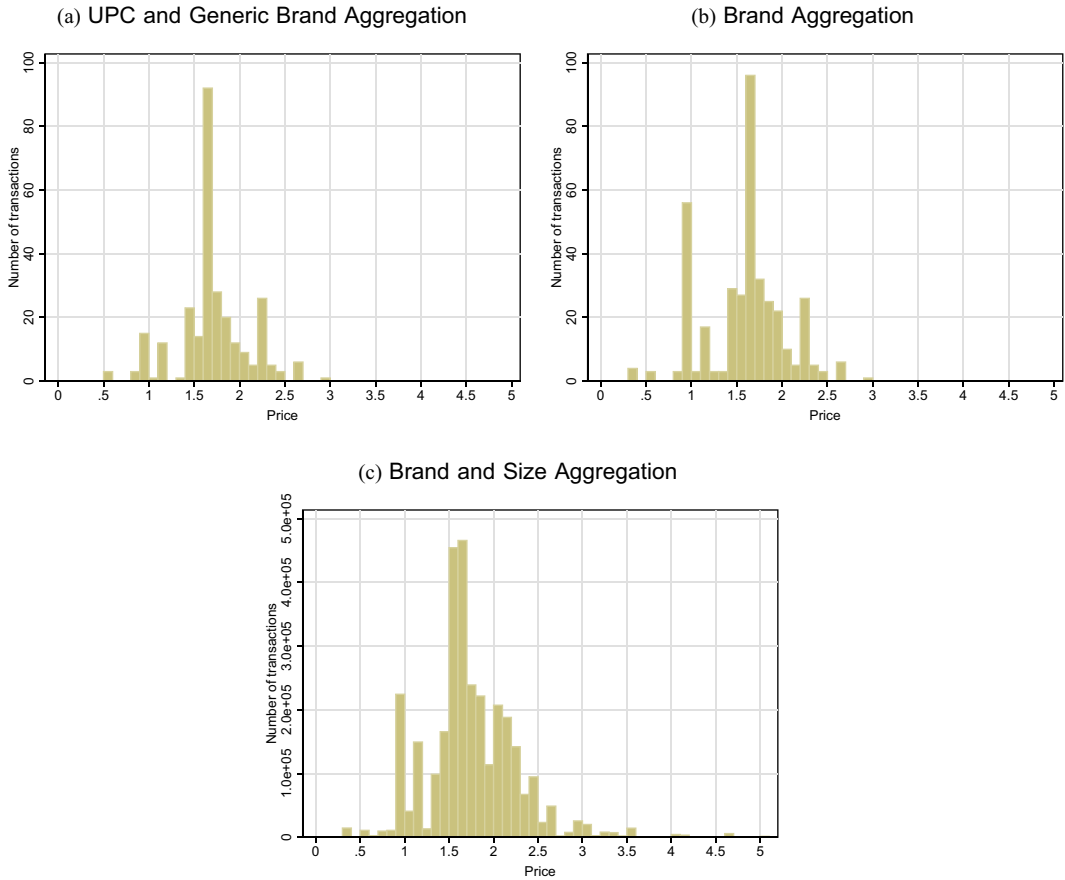
Second, in order to limit the possibility that the three broader definitions of a good introduce some misclassification error (perhaps due to unobserved product features or errors in the size information), we drop all goods in a given market m and quarter t for which the amount of price dispersion is unreasonably large. We set the threshold at a coefficient of variation greater than 1. In our experience, misclassification occurs infrequently and generates very large measured price dispersion.

Third, in order to avoid the possible influence of a small number of outliers on measured price dispersion, we drop all transactions in which the recorded price is zero. Although some transactions with zero price are indeed valid (e.g., store freebies or manufacturer coupons), the KNCP data do not distinguish between valid zero price transactions and transactions for which price data was not available. Similarly, we exclude all transactions involving a coupon whose value is more than 90% of the gross price. We have experimented with setting this threshold at 50% and with excluding transactions with coupons altogether.

3. PRICE DISPERSION

In this section, we document the extent of price dispersion and how it varies across goods, markets, and time. We then describe and implement a decomposition of the variance of prices into store, good, and transaction components, discussing the implications of our decomposition for existing theories of price dispersion. For the sake of concreteness, we first illustrate price dispersion for one particular good in one particular market and quarter.

3.1. *An Example: Heinz Ketchup.* As an illustrative example, we take 36-oz plastic bottles of Heinz ketchup purchased in Minneapolis in 2007:Q1. Heinz ketchup is part of the Catsup



NOTES: Figures show distribution of transaction prices for ketchup in Minneapolis in 2007:Q1. Panel (a) shows prices for 36-oz bottles of Heinz brand ketchup, in accordance with the UPC and Generic Brand Aggregation definitions of a good. Panel (b) shows prices for 36-oz bottles of ketchup from all brands, in accordance with the Brand Aggregation definition of a good. Panel (c) shows prices per 36 ounces of ketchup for bottles of ketchup from all brands in all sizes, in accordance with the Brand and Size Aggregation definition of a good. Transactions include those at all stores in Minneapolis, including stores without a unique identifier.

FIGURE 1

DISTRIBUTION OF PRICES FOR A 36-OZ BOTTLE OF HEINZ KETCHUP

product module, which is one of 34 product modules in the Condiments, Gravies, and Sauces product group, which is one of 38 product groups in the Dry Grocery department.

In Minneapolis in 2007:Q1, we observe 279 transactions for 36-oz plastic bottles of Heinz ketchup. Figure 1(a) shows the distribution of prices across these 279 transactions. The lowest price in the distribution is 50 cents. The highest price in the distribution is \$2.99, which is six times larger than the lowest price. The median price is \$1.66, which is more than three times larger than the lowest price and roughly half of the highest price. The coefficient of variation of prices is 23%. Since all 36-oz plastic bottles of Heinz ketchup have the same UPC, they represent the same good according to the UPC definition. Since there are no other products with the same product features, size, and brand, 36-oz plastic bottles of Heinz ketchup also represent a single good according to the Generic Brand Aggregation definition. Hence, the price distribution in Figure 1(a) applies to the two narrowest definitions of a good.

According to the Brand Aggregation definition of a good, 36-oz plastic bottles of Heinz ketchup constitute a good together with 36-oz plastic bottles of ketchup from other brands. In Minneapolis in 2007:Q1, these are Frank's, Hunt's, and nine generic brands. Figure 1(b) shows

TABLE 2
AVERAGE STATISTICS OF PRICE DISTRIBUTION

	Aggregation UPC (1)	Generic Brand Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Standard deviation	0.19	0.21	0.25	0.36
90–10 ratio	1.72	1.79	2.04	2.61
90–50 ratio	1.26	1.29	1.38	1.55
50–10 ratio	1.35	1.38	1.46	1.64

NOTES: All statistics are expenditure weighted across goods, markets, and quarters. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1.

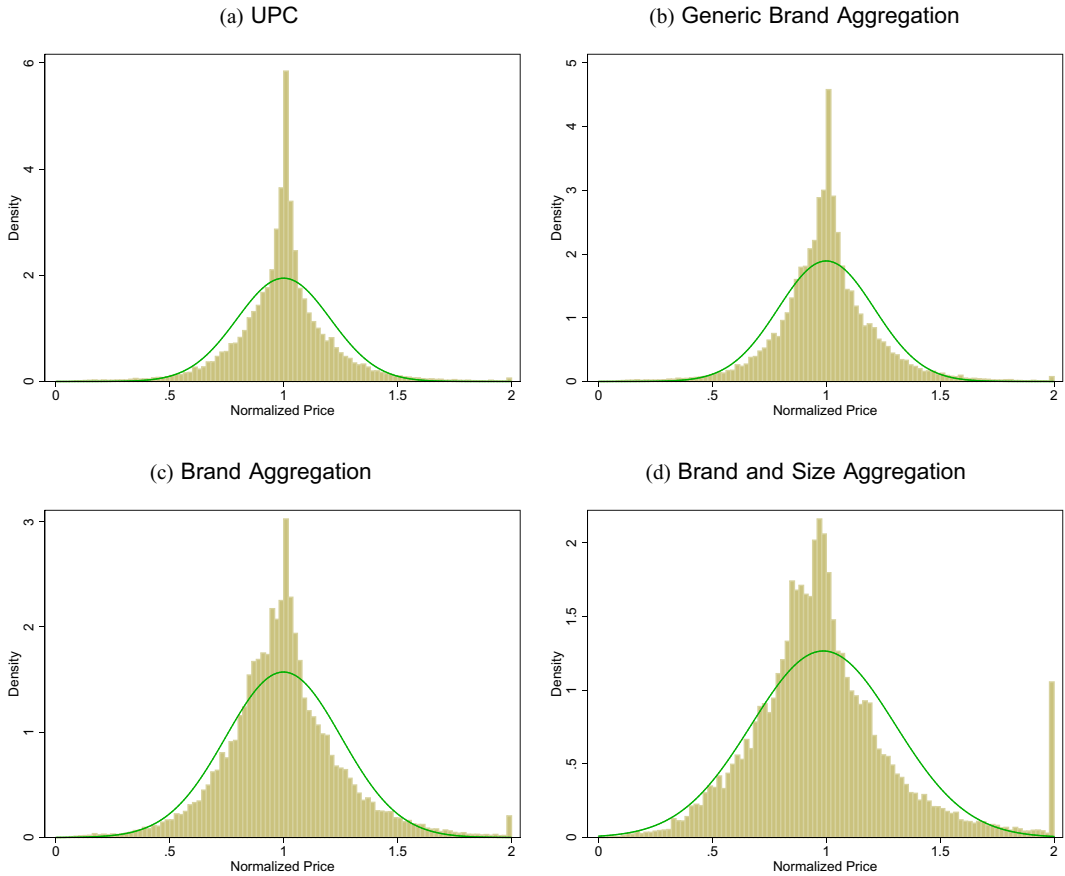
the distribution of prices for 36-oz plastic bottles of ketchup of these 12 brands. Compared to the distribution in 1(a) there is more price dispersion: the coefficient of variation is 27%, the lowest price is 39 cents, and the highest price is \$2.99.

According to the Brand and Size Aggregation definition of a good, 36-oz plastic bottles of Heinz ketchup constitute a good together with other size plastic bottles of ketchup from both Heinz and other brands. In Minneapolis in 2007:Q1, we observe transactions for 53 such ketchup bottles spanning 12 brands and 13 sizes. Figure 1(c) shows the distribution of prices for bottles of ketchup of all brands and sizes. Compared to the distribution in Figure 1(b) there is even more price dispersion: the coefficient of variation is 29%, the lowest price is 1.08 cents per ounce (equivalent to 39 cents for a 36-oz bottle) and the highest price is 13.96 cents per ounce (equivalent to \$5.03 for a 36-oz bottle).

3.2. Measuring Price Dispersion. The amount of variation in the prices of ketchup bottles is rather typical of the goods in our data set. We compute statistics of the normalized transaction price distribution separately for every good j in every market m and quarter t . The expenditure-weighted averages of these statistics (across goods, markets, and quarters) are reported in Table 2. When goods are defined according to their UPC, the average standard deviation of prices is 19%, the average 90–10 percentile ratio is 1.72, the average 90–50 percentile ratio is 1.26, and the average 50–10 percentile ratio is 1.35. These numbers reveal that there is substantial dispersion in the price at which the same good is sold in a given market and quarter. Moreover, the similarity of the average 90–50 ratio and the average 50–10 ratio implies that price distributions are roughly symmetric. Hence, the price dispersion that we observe is not driven by a small number of very low-priced transactions—for example, due to temporary sales.

As we apply broader definitions of a good, the amount of price dispersion increases. The second, third, and fourth columns of Table 2 show that the average standard deviation of prices increases to 21% when we aggregate products with different generic brands, to 25% when we aggregate products with different name brands, and to 36% when we aggregate products with different brands and sizes. Similarly, the average 90–10 percentile ratio increases to 1.79 when we aggregate products with different generic brands, to 2.04 when we aggregate products with different name brands, and rises to 2.61 when we aggregate products with different brands and sizes.

In Figure 2, we illustrate the average price distribution for each definition of a good, by plotting the distribution of normalized prices across all markets, goods, and quarters. The figure reveals several features of the shape of a ‘typical’ price distribution. First, for all definitions of a good, the price distribution has a unique mode that is very close to its mean. Second, for all definitions of a good, the price distribution is very close to symmetric. Third, for all definitions of a good, the price distribution has more mass around the mean and has thicker tails than a Normal distribution with the same mean and variance; in other words, the typical price distribution is leptokurtic.



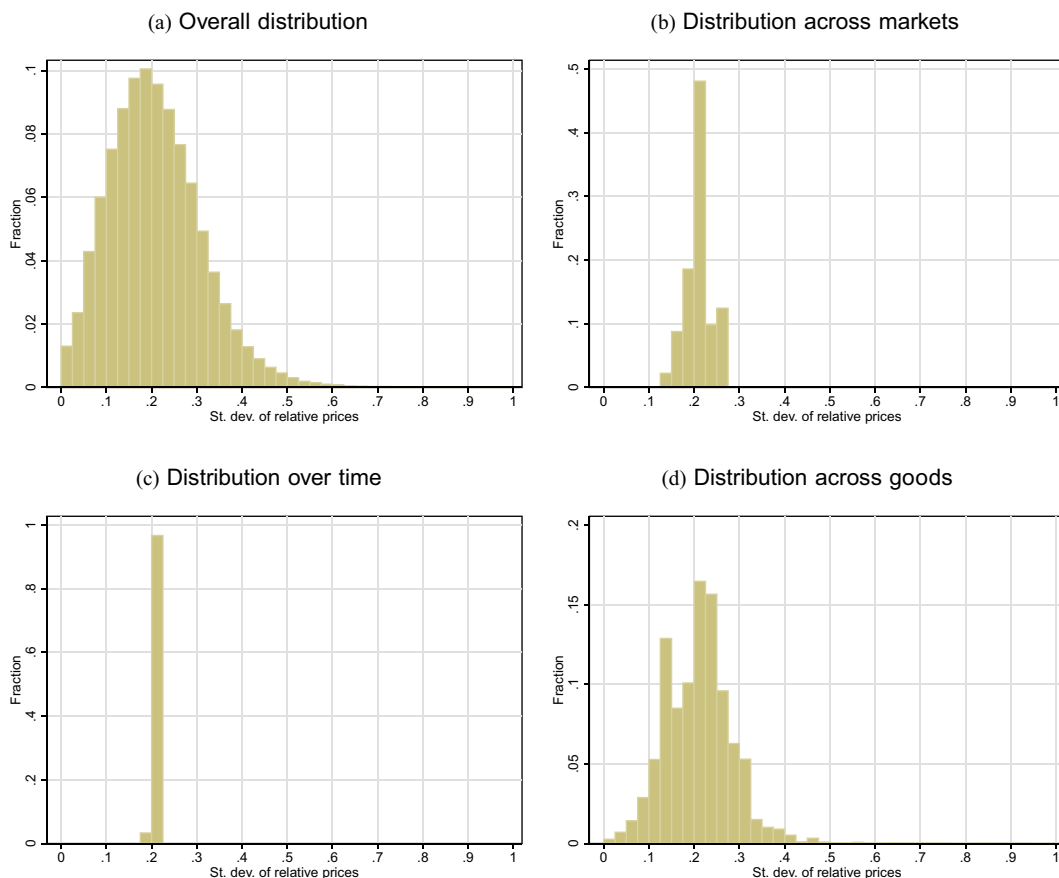
NOTES: Figures plot average distribution of normalized prices across all goods, markets, and time periods for different definitions of a good. All distributions are expenditure weighted. Overlaid curve is the density of a normal distribution with the same mean and variance. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1.

FIGURE 2

AVERAGE PRICE DISTRIBUTION

We saw in Table 2 that, on average, there is substantial dispersion in the prices paid for the same good in a particular market and quarter. Figure 3(a) plots the distribution of the standard deviation of prices across different goods, markets, and quarters and reveals that there is also significant heterogeneity in the extent of price dispersion. For example, while the average standard deviation of prices is 21% for the Generic Brand Aggregation definition, there are approximately 10% of goods/markets/quarters where the standard deviation of prices is lower than 10% and there are approximately 10% of goods/markets/quarters where the standard deviation of prices is greater than 35%.

The heterogeneity in price dispersion across goods, markets, and quarters is primarily due to the fact that price dispersion differs systematically across goods, not because the amount of price dispersion varies substantially across markets or over time. To illustrate this, the remaining panels of Figure 3 each show a histogram of the standard deviation of relative prices across only one of these dimensions. The histograms are plotted on the same scale to aid comparison. They show that there is little variation in the average over goods and quarters of the standard deviation of prices in different markets (Figure 3(b)). There is also little variation in the average over goods and markets of the standard deviation of prices in different quarters (Figure 3(c)). In



NOTES: Figures refer to generic brand aggregation definition of a good. All distributions are expenditure weighted. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1.

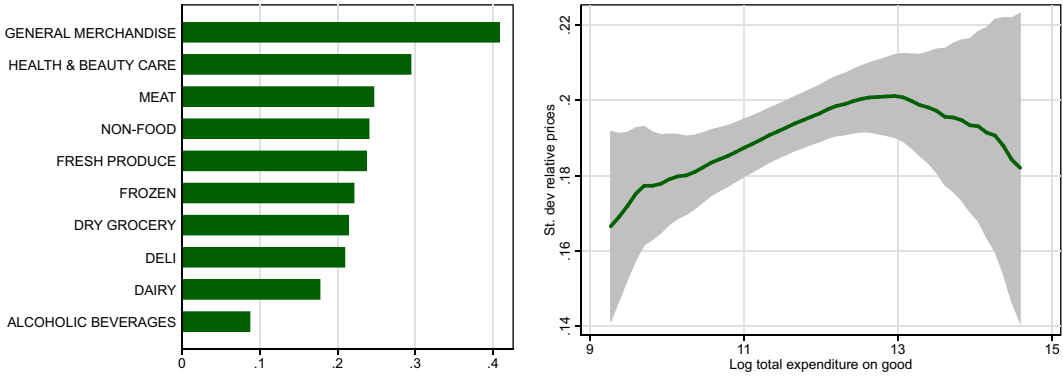
FIGURE 3

DISTRIBUTION OF STANDARD DEVIATION OF PRICES ACROSS GOODS, MARKETS, AND TIME

contrast, there is substantial variation in the average over markets and quarters of the standard deviation of prices for different goods (Figure 3(d)).

It is then natural to ask which types of goods feature more price dispersion. Figure 4(a) plots the average standard deviation of prices for goods in the 10 different departments that Nielsen uses to classify products. The two departments with the highest standard deviation of prices are General Merchandise (41%) and Health and Beauty Care (29%). The two departments with the lowest standard deviation of prices are Alcoholic Beverages (8%) and Dairy (17%). Figure 4(b) plots the average standard deviation of the price of a good against the log of total market expenditure for the good for Minneapolis in 2007:Q1. The figure shows a distinct hump-shaped relationship between price dispersion and the expenditure share of a good. Other markets and time periods show a similar relationship. We return to this relationship in Subsection 3.4 in our discussion of alternative theories of price dispersion.

3.3. Deconstructing Price Dispersion. Consider two of the 36-oz bottles of Heinz ketchup that were purchased in Minneapolis in 2007:Q1. The difference in prices between these two bottles of ketchup may have originated from three conceptually different sources. First, there may be differences across stores in the average price of goods. If so, the two bottles of ketchup may have sold at different prices because one was purchased at a relatively cheap store and



NOTES: (a) Left panel shows expenditure-weighted average standard deviation of prices, across all markets and time periods, for goods in each department. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1. (b) Right panel shows nonparametric estimate of relationship between standard deviation of prices and log total market expenditure across goods. Estimation is by local linear regression with Epanechnikov kernel, with bandwidth = 0.54. Shaded area is 95% confidence interval estimated with bandwidth = 0.81. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1 in Minneapolis in 2007:Q1.

FIGURE 4

AVERAGE PRICE DISPERSION FOR DIFFERENT GOODS

one at a relatively expensive store. Second, there may be heterogeneity in the relative prices of different goods within stores. If so, the two bottles of ketchup may have sold at different prices because they were purchased at two different, but equally expensive, stores. This could happen if one of those stores charges a relatively low price for ketchup (offset by a relatively high price for some other products), and the other store charges a relatively high price for ketchup (offset by a relatively low price for some other products). Third, the two bottles of ketchup may have been purchased at exactly the same store, but at different prices. This could happen either because one of the bottles was purchased on a day during the quarter when these bottles of ketchup were on a temporary sale or because one of the bottles was purchased with a coupon. The KNCP data contain information on the transaction prices for different goods at the same store and on the transaction prices for the same good at different stores in a given market and quarter. This allows us to decompose the observed price dispersion into these three sources of variation.

To formalize our decomposition of price dispersion, we require some additional notation. First, let $\mu_{j,m,t}$ denote the quantity-weighted average of the normalized prices for good j in market m and quarter t . Formally, $\mu_{j,m,t}$ is defined as

$$(3) \quad \mu_{j,m,t} = \frac{\sum_{k \in m \cap t} p_{j,k} q_{j,k}}{\sum_{k \in m \cap t} q_{j,k}}.$$

Note that $\mu_{j,m,t}$ is always equal to 1 because the prices $p_{j,k}$ are defined as the actual prices relative to the quantity-weighted average price of good j in market m and quarter t .

Second, let $\mu_{j,s,t}$ denote the quantity-weighted average of the normalized prices for good j at store s in quarter t . Formally, $\mu_{j,s,t}$ is defined as

$$(4) \quad \mu_{j,s,t} = \frac{\sum_{k \in s \cap t} p_{j,k} q_{j,k}}{\sum_{k \in s \cap t} q_{j,k}}.$$

Third, we denote the revenue-weighted average of the price of the different goods sold at store s by $\mu_{s,t}$. Formally, $\mu_{s,t}$ is defined as

$$(5) \quad \mu_{s,t} = \frac{\sum_j \mu_{j,s,t} R_{j,s,t}}{\sum_j R_{j,s,t}},$$

where the revenue, $R_{j,s,t}$, is defined by

$$(6) \quad R_{j,s,t} = \sum_{k \in s \cap t} P_{j,k} q_{j,k}.$$

Using the above definitions, we can decompose any normalized price as

$$(7) \quad p_{j,k} = \mu_{j,m,t} + (\mu_{s,t} - \mu_{j,m,t}) + (\mu_{j,s,t} - \mu_{s,t}) + (p_{j,k} - \mu_{j,s,t}),$$

where it is implicit that the market m and store s refer to $m(i(k))$ and $s(k)$, respectively, i.e., the market and store corresponding to shopping trip k .

The first term on the right-hand side of (7) is the average price of good j in market m and quarter t . The second term on the right-hand side of (7) is the difference between the average price of store s and the average price of good j . This term measures the component of the price $p_{j,k}$ that is due to the expensiveness of the store where the good was purchased. We refer to this as the *store component* of the price. The third term on the right-hand side of (7) is the difference between the average price of good j at store s and the average price of store s . This term measures the component of the price $p_{j,k}$ that is due to the expensiveness of good j relative to other goods at store s . We refer to this as the *store-specific good component* of the price. The last term on the right-hand side of (7) is the difference between the price $p_{j,k}$ and the average price of good j at store s . This term measures the component of the price $p_{j,k}$ that is due to the expensiveness of the particular transaction relative to the average expensiveness of good j at store s . We refer to this as the *transaction component* of the price.

The decomposition in (7) is exact. Therefore, the variance of normalized transaction prices for good j in market m and quarter t must equal the sum of the variances of the store component, the store-specific good component, and the transaction component, and the covariances between these three components. In Appendix A, we show that the covariance between the store and transaction components, as well as the covariance between the store-specific good and transaction components always equals zero. Empirically, the covariance between the store and store-specific good components is not zero, but it is generally a negligible component of the overall variance of prices.¹²

Although the variance decomposition in (7) is simply a description of the data and as such is always valid, for matters of interpretation it is important to keep in mind that the average price

¹² We refer to quantity-weighted variances. See Appendix 5A for a full derivation of the variance decomposition and formulas for the variances of each term.

TABLE 3
DECOMPOSITION OF THE VARIANCE OF PRICES

	Substitution UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
A: Restrict to stores/goods with at least five transactions				
Store-specific component	6%	7%	8%	8%
Store-good specific component	24%	30%	38%	40%
Transaction component	70%	62%	53%	53%
Covar: store, store-good	-1%	0%	1%	-2%
B: Common variance of the transaction component across goods				
Store-specific component	11%	11%	13%	15%
Store-good component	38%	41%	43%	31%
Transaction component	49%	47%	46%	64%
Covar: store, store-good	2%	1%	-2%	-10%

NOTES: All statistics are expenditure weighted across goods, markets and quarters. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1. See text for details of decomposition. May not add to 100% due to rounding.

of goods at store s , $\mu_{s,t}$ and the average price of good j at store s , $\mu_{j,s,t}$ are sample averages in the KNCP data, instead of averages in the universe of transactions at store s . Since the number of observed transactions at a given store is very large, the sample average price of goods at store s is likely to be a precise estimate of the average price of goods at store s in the universe of transactions. However, there are several store/good combinations for which the number of observed transactions is small. For these store/good combinations, the sample average price of the good at that store is not likely to be a precise estimate of the average in the universe of transactions.

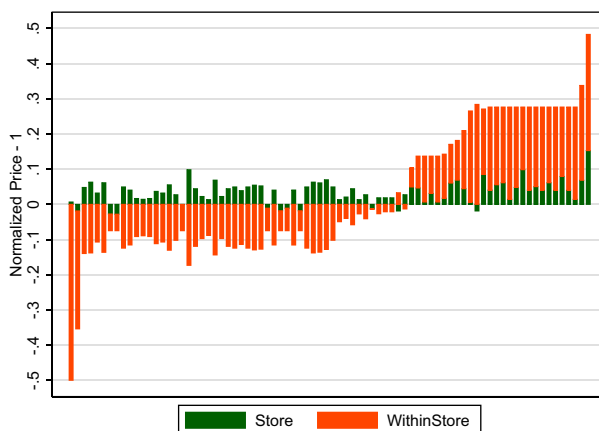
We employ two alternative approaches to deal with the scarcity of data for some store/good combinations. The first approach is to restrict attention to stores and goods with a minimum number of transactions, which we set at 5.¹³ The variance decomposition under this restriction is reported in Panel A of Table 3. For the Generic Brand Aggregation definition of a good (column (2)), the store component accounts for approximately 7% of the variance of prices, the store-specific good component accounts for approximately 30% of the variance of prices, and the transaction component accounts for approximately 62% of the variance of prices. The covariance between the store and store-specific good components accounts for less than 1% of the variance of prices.

The second approach that we use to overcome the scarcity of data for some store/good combinations is to impose structure on the variance of the transaction component of prices. In particular, we assume that the transaction component of the price for good j at store s in market m in quarter t is drawn from a distribution with mean zero and variance $\sigma_{j,m,t}^2$, where $\sigma_{j,m,t}^2$ is modeled as the sum of a quarter-specific market effect and a quarter-specific good department effect. Given this assumption, we can estimate the variance of the transaction component by pooling data across different stores and goods in the same market and quarter. We then recover the variance of the store-specific good component using the fact that

$$(8) \quad \text{var}(\mu_{j,s,t} - \mu_{s,t}) = \text{var}(p_{j,k} - \mu_{s,t}) - \sigma_{j,m,t}^2.$$

The resulting decomposition is reported in Panel B of Table 3. For the Generic Brand Aggregation definition of a good (column (2)), the store component accounts for approximately

¹³ Note that since we compute the decomposition separately for each good, market, and quarter and report the expenditure-weighted averages, the decomposition without any restrictions already weights more heavily the goods, markets, and quarters for which we observe more transactions. In addition, in all of these analyses we have already restricted attention to goods, markets, and quarters with at least 25 transactions (spread across all stores). We have also repeated the decomposition using a minimum of 10 transactions and have found similar qualitative results.



NOTES: Figure shows 107 transactions for 36-oz plastic bottles of Heinz ketchup in Minneapolis in 2007:Q1. Includes only transactions at stores with a unique store identifier.

FIGURE 5


PRICE DECOMPOSITION FOR KETCHUP IN MINNEAPOLIS, 2007:Q1

11% of the variance of prices, the store-specific good component accounts for approximately 42% of the variance of prices, and the transaction component accounts for approximately 46% of the variance of prices. The covariance between the store and store-specific good components accounts for 1% of the variance of prices. We find that, consistent and in keeping with the decompositions in Panel A, the store component is relatively unimportant for explaining price dispersion whereas the transaction component is somewhat more important than the store-specific good component.

Thus, the results from our decomposition of the variance of prices can be summarized as follows: The store component explains approximately 10% of the overall variance of prices whereas the sum of the store-specific good component and the transaction component, which we refer to as the *within-store component* of the price, accounts for approximately 90% of the overall variance of prices. The contribution of the store-specific good component varies between 25% and 45%, depending on both how we define a good and on the assumptions we make about the variance of the transaction component to overcome the scarcity of data for some store/good combinations. Conversely, the contribution of the transaction component varies between 45% and 65%.

Perhaps the most striking finding from our decomposition is that the store component explains very little of the overall variance of prices. To illustrate this result more concretely, Figure 5 decomposes the price of every transaction for 36-oz plastic bottles of Heinz ketchup in Minneapolis in 2007:Q1. The transactions are ordered from cheapest to most expensive. For each transaction, the darker bar measures the store component of the price, $\mu_{s,t} - \mu_{j,m,t}$, whereas the lighter bar measures the within-store component of the price, $p_{j,k} - \mu_{s,t}$. From Figure 5, it is evident that the difference in the expensiveness of the stores where each bottle of ketchup was purchased accounts for very little of the overall dispersion in prices. Instead, the dispersion in prices is mostly accounted for by the fact that the same bottle of ketchup is sold at vastly different prices either at different stores that are equally expensive or at the very same store on different days or in different transactions on the same day.

To further demonstrate the relative importance of the within-store component of prices and the relative unimportance of the store component of prices, Figure 6 shows a scatter plot of all the transactions that took place in the first quarter of 2007. Each point in the plot represents a transaction. The horizontal axis measures the store component of the price at which the transaction was executed, $\mu_{s,t} - \mu_{j,m,t}$. The vertical axis measures the normalized price for the transaction relative to the average normalized price for that good, $p_{j,k} - \mu_{j,m,t}$. If all price



NOTES: Figure plots the normalized price minus $\mu_{j,m,t}$ against the store component of the price for all transactions involving goods that are purchased at least five stores in all markets in 2007:Q1. Approximately 0.3% of transactions are truncated so as to stay in the range $[-1,2]$.

FIGURE 6

STORE COMPONENT VERSUS NORMALIZED PRICE

dispersion was due to differences in the expensiveness of stores, all points would lie on the 45° line. In contrast, if all price dispersion was due to differences in prices at equally expensive stores, all points would lie on a vertical line passing through the origin. The fact that the points are much more dispersed around the 45° line than they are along the vertical line passing through the origin clearly illustrates the fact that price dispersion arises because, at stores that are equally expensive, there are goods that are sold at very high prices (relative to the market average) and there are goods that are sold at very low prices (relative to the market average). In turn, this implies that the set of goods that are sold at low and high prices must vary across stores that are equally expensive.

3.4. Price Dispersion: Evidence and Theory. Our decomposition of the variance of prices for identical goods enables us to assess the relative importance of four popular explanations for the existence of price dispersion in product markets: (i) amenities; (ii) heterogeneous costs; (iii) intertemporal price discrimination; and (iv) search frictions.¹⁴

3.4.1. Amenities. According to the amenity theory of price dispersion, product markets are frictionless in the sense that buyers know all sellers' prices and are free to purchase from any seller. Yet identical goods sell at different prices because they are bundled with different amenities in different transactions. For example, an identical bottle of ketchup will be sold at a higher price at a store with convenient parking or a high cashier–customer ratio and at a lower price at a store without convenient parking or with a low cashier–customer ratio.

Since amenities are typically a feature of a store (such as the location of the store, the quantity and quality of the sales staff, or the presence of convenient parking), high-amenity stores will

¹⁴ In this section, we only review theories of price dispersion among firms selling a single product, as this is the assumption commonly made in the literature. There are only a handful of theoretical papers studying the optimal pricing of firms selling multiple products. McAfee (1995) proposes a search-theoretic model of pricing for multiproduct firms in the spirit of Burdett and Judd (1983). He finds conditions for the existence of an equilibrium in which the distribution of prices for each good is the same as in a model where each firm only sells one good. Zhou (2014) considers a search-theoretic model of pricing of multiproduct firms in the spirit of Wolinsky (1986). However, the focus is on equilibria without price dispersion across firms. McAfee et al. (1989) and Chu et al. (2011) consider the problem of a monopolist selling multiple goods. However, the subject of these papers is the optimality of bundling instead of price dispersion.

TABLE 4
TRANSITION MATRIX OF STORE PRICES

Quartile	Quartile			
	Bottom	Second	Third	Top
Bottom	42%	24%	15%	19%
Second	20%	37%	27%	16%
Third	12%	26%	39%	23%
Top	18%	18%	27%	37%

NOTES: Table shows the fraction of stores transiting across quartiles of the distribution of estimated store prices.

charge relatively high prices for all goods, whereas low-amenity stores will charge relatively low prices for all goods. Moreover, if the extent to which customers enjoy the amenities of a store is roughly proportional to the value of the goods that they purchase at that store, the ratio between the price of a good at a high-amenity store and a low-amenity store should be approximately the same for all goods. Under this assumption, the variance of prices caused by differences in amenities should be accounted for mostly by the store component of prices. Table 3 shows that the variance of the store component of prices accounts for at most 15% of the overall variance of prices. Thus, differences in amenities are not likely to be a quantitatively important source of price dispersion. This is the same conclusion that Sorensen (2000) reaches by analyzing the variation in retail prices for prescription drugs across 20 pharmacies in two towns in upstate New York.

Table 4 shows that differences in amenities are not even likely to account for all of the small variance of the store component of prices. The table reports the annual transition rates across quartiles of the distribution of estimated store prices. Specifically, the entry in row i and column j is the fraction of stores in the i th quartile of the distribution of store prices in quarter t , $\mu_{s,t}$, whose average price one year later $\mu_{s,t+4}$ is in the j th quartile of the store price distribution in quarter $t + 4$. Since most amenities are associated with persistent features of a store (location, parking, staff quality), if the store component mostly reflects amenities, then the average price of the store should be persistent over time. Table 4 shows that this is not the case. Indeed, 19% of the stores whose price is in the bottom quartile in a given quarter are in the top quartile of the store price distribution one year later. Similarly, 18% of the stores whose price is in the top quartile in a given quarter are in the bottom quartile of the store price distribution one year later. Overall, the annual rank autocorrelation for stores prices $\mu_{s,t}$ is only 25%.

The discussion above rests on the assumption that the extent to which customers enjoy the amenities of a store is proportional to the value of the goods that they purchase at the store. If this assumption does not hold, then the ratio between the price of a good at a high-amenity store and a low-amenity store may vary across goods. In this case, the price dispersion that arises due to differences in amenities would be accounted for by both the variance of the store component and the variance of the store-specific good component. However, even if this were the case, high-amenity stores would charge higher prices than low-amenity stores for all goods. But Figure 6 shows that this is far from true in our data. Indeed, stores that are more expensive than average (they are toward the right on the horizontal axis) sell many goods at below-average prices (goods that are below zero on the vertical axis) and, conversely, stores that are less expensive than average sell many goods at above-average prices.

3.4.2. *Heterogeneous costs.* According to this theory, identical goods are traded at different prices because they are sold by local monopolists who face different marginal costs (see, e.g., Golosov and Lucas, 2007). For example, an identical bottle of ketchup will be sold at a higher price by a monopolist who faces a higher wholesale price or a higher cost of storage than by a monopolist who faces a lower wholesale price or a lower cost of storage.

It seems natural to think that differences in marginal costs arise mostly because different sellers pay different prices in the wholesale market. Since differences in wholesale prices across sellers may vary across goods, heterogeneity in marginal costs could account for both the variance of the store component of prices and the variance of the store-specific good component of prices. This amounts to a sizable fraction of the overall variance of prices. However, such a conclusion would overstate the importance of heterogeneity in marginal costs because sellers that are part of the same retail chain typically bargain collectively with wholesalers and so should face the same wholesale prices and thus similar marginal costs.¹⁵ Therefore, heterogeneity in marginal costs can explain, at most, the variance of the components of prices that are common to all sellers in the same retail chain.

The KNCP data allow us to decompose each transaction price into a retailer component, a retailer-specific store component, a retailer-specific good component, and a residual component. Formally, we decompose any normalized price $p_{j,k}$ as

$$(9) \quad p_{j,k} = \mu_{j,m,t} + (\mu_{r,m,t} - \mu_{j,m,t}) + (\mu_{s,t} - \mu_{r,m,t}) \\ + (\mu_{j,r,m,t} - \mu_{r,m,t}) + (p_{j,k} - \mu_{s,t} - \mu_{j,r,m,t} + \mu_{r,m,t}),$$

where $\mu_{j,m,t}$ is the average price of good j in market m in quarter t , $\mu_{s,t}$ is the average price of goods at the store $s(k)$, $\mu_{r,m}$ is the average price of goods at stores in market $m(i(k))$ that belong to the retail chain $r(k)$, and $\mu_{j,r,m,t}$ is the average price of good j at stores in market $m(i(k))$ that belong to retail chain $r(k)$.

The first term on the right-hand side of (9) is the market average price of good j . The second term on the right-hand side of (9) is the difference between the average price of retail chain r in market m and the average price of good j in market m . We refer to this as the *retailer component* of the price. The third term on the right-hand side of (9) is the difference between the average price of store s in retail chain r and the average price of retail chain r in market m . We refer to this as the *retailer-specific store component* of the price. The fourth term on the right-hand side of (9) is the difference between the average price of good j at retail chain r in market m and the average price of retail chain r in market m . We refer to this as the *retailer-specific good component* of the price. The fifth term on the right-hand side of (9) is a residual component. The components of the price that are common to all retailers in the same chain are the retailer component and the retailer-specific good component.

Using the identity in (9), we decompose the variance of prices into a retailer component, a retailer-specific store component, a retailer-specific good component, and a residual component and report the results in Table 5.¹⁶ The decomposition shows that the retailer component accounts for less than 3% of the overall variance of prices. A comparison between Table 5 and Panel B of Table 3 reveals that the fraction of price dispersion that is accounted for by the retailer component is much smaller than what is accounted for by the store component of prices. This implies that the average price of a store depends less on the retail chain to which the store belongs and more on the characteristics that are idiosyncratic to the store (e.g., its location). Table 5 also shows that the variance of the retailer-specific good component accounts for 7%–14% of the overall variance of prices, depending on the definition of a good. Again, the fraction of price dispersion that is accounted for by the retailer-specific good component is much smaller than the fraction that is accounted for by the store-specific good component of prices. This implies that the relative price of a particular good at a store also depends less on the retail chain to which the store belongs and more on characteristics that are idiosyncratic

¹⁵ For example, in a study of a supermarket chain in the Basque region of Spain, Aguirregabiria (1999) finds that all wholesale purchases are made by one central store. Moreover, he reports that, according to the purchasing managers of this chain, wholesale prices are generally the same for all the supermarket chains in the region.

¹⁶ To deal with the scarcity of data for some store/good combinations, the decomposition in Table 5 is computed using only data from the store/good combinations for which we observe at least five transactions, as described in Subsection 3.3.

TABLE 5
DECOMPOSITION OF VARIANCES OF PRICES BY RETAILER

	Substitution UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Retailer specific component	1%	3%	3%	2%
Retailer-store component	5%	5%	6%	6%
Retailer-good component	7%	13%	14%	10%
Residual component	90%	82%	80%	85%
Covar: retailer, retailer-store	0%	0%	0%	0%
Covar: retailer-good, residual	-2%	-3%	-3%	-1%
Covar: store, store-good	-1%	0%	1%	-2%

NOTES: All statistics are expenditure weighted across goods, markets, and quarters. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1. Sample is further restricted to good/store/quarter combinations with at least five transactions. See text for details of decomposition. May not add to 100% due to rounding.

to the store. In general, very little of the overall variance in prices can be explained by the components of prices that are common to all stores in the same retail chain. Indeed, Table 5 shows that the sum of the variances of the retailer component and the retailer-specific good component together account for 8%–17% of the overall variation in prices for identical goods. Thus, our results suggest that heterogeneity in marginal costs is not a quantitatively important explanation for price dispersion.

3.4.3. Intertemporal price discrimination. According to intertemporal price discrimination theories, price dispersion arises because local monopolists change their prices over time in order to discriminate between different types of buyers (see Conlisk et al., 1984; Sobel, 1984; Albrecht et al., 2013). For the sake of illustration, consider a market populated by several monopolists, each selling the same good and facing a constant inflow of heterogeneous buyers. The first type of buyer has a relatively low valuation for the good and has a relatively high willingness to substitute consumption over time. The second type of buyer has a relatively high valuation for the good and is less willing to substitute consumption over time. In this environment, a monopolist finds it optimal to charge a relatively high price for some interval of time. Although the monopolist charges this relatively high price, high-valuation buyers purchase the good and low-valuation buyers wait. When the number of low-valuation buyers who are waiting becomes sufficiently large, the monopolist finds it optimal to lower the price for a short interval of time. At this lower price, all the low-valuation buyers purchase the good and exit the market, after which the pricing cycle starts over. Indeed, a number of authors have found evidence for this type of pricing behavior using store-level data (see Klenow and Malin, 2010, for an overview).¹⁷

If pricing cycles are short relative to the unit of time over which we measure price dispersion (a quarter), the price dispersion that originates from intertemporal price discrimination should be accounted for mostly by the variance in the transaction component of prices. Table 3 reveals that the variance of the transaction component of prices represents somewhere between 45% and 65% of the overall variance of prices. Thus, intertemporal price discrimination could be the source of a significant fraction of the observed variation in prices for identical goods.

If pricing cycles are long relative to a quarter, then price dispersion that originates from intertemporal price discrimination would be accounted for by both the variance of the

¹⁷ Intertemporal price discrimination is not the only theory that can account for pricing cycles. For example, Aguirregabiria (1999) argues that temporary sales may be related to inventory management. When sellers restock their inventories they face a low probability of a stockout and hence a higher effective elasticity of demand. This leads to lower prices. As sellers run down their inventories, the stockout probability increases, the effective elasticity of demand falls, and prices increase. The implications for price dispersion of this stockout theory of pricing cycles are the same as the implications of the intertemporal price discrimination theory.

store-specific good component of prices and the variance of the transaction component of prices. In this case, intertemporal price discrimination could explain up to 90% of the overall variance of prices. However, we find evidence suggesting that pricing cycles are short. Indeed, if pricing cycles were long relative to the unit of time over which we measure price dispersion, an increase in the unit of time should move some of the variance of prices from the store-specific good component to the transaction component. Yet, when we increase the unit of time from one quarter to one year, we see very little change in the size of the transaction component of prices relative to the store-specific good component.

3.4.4. Search frictions. The presence of search frictions in the product market can simultaneously explain why buyers do not arbitrage away price differences and why sellers choose to charge different prices (see Butters, 1977; Varian, 1980; Burdett and Judd, 1983). For the sake of illustration, consider a market populated by a large number of buyers and sellers. Due to search frictions, an individual buyer can only purchase from some subset of sellers at any given time. In particular, some buyers are able to purchase from multiple sellers, whereas other buyers are only able to purchase from one seller. In this environment, if all sellers charged the same price, an individual seller could increase its profits by posting a slightly lower price and selling to the buyers who can purchase from someone else as well as to the buyers who cannot purchase from any other sellers. Thus, sellers must randomize with respect to the price they charge and price dispersion arises.

Depending on the process of randomization followed by the sellers, the price dispersion due to search frictions could show up in the variance of any of the three components of prices. If the sellers' randomization process is perfectly correlated across goods and times, the price dispersion due to search frictions would show up in the variance of the store component of prices. If the sellers' randomization process is independent across goods and correlated over time, the price dispersion due to search frictions would show up in variance of the store-specific good component of prices. If the sellers' randomization process is independent across goods and over time, the price dispersion due to search frictions would show up in the variance of the transaction component of prices. Indeed, given the appropriate pattern of randomization, the price dispersion due to search frictions could show up in the variance of the three components of prices in the same proportions as we observe in the KNCP data.

However, it is unlikely that the variance of the transaction component of prices is due to search frictions because, in the presence of any cost of price adjustment, sellers would refrain from resetting their prices very often. Thus, search frictions are likely to be responsible for, at most, the variance of the store component of prices and the variance of the store-specific good component of prices, which, according to the decompositions in Table 3, represent somewhere between 35% and 55% of the overall variance of prices.

The variation in price dispersion across different types of goods is also consistent with the view that search frictions are an important component of overall price dispersion. Figure 4(b) shows a nonmonotonic, hump-shaped relationship between the average standard deviation of the price of a good and the share of household expenditures devoted to that good. This is precisely the relationship implied by the search-theoretic models of Butters (1977), Varian (1980), and Burdett and Judd (1983), assuming that households have a stronger incentive to seek multiple sellers for goods that account for a larger fraction of their expenditures. For goods that account for a small share of household expenditures, the fraction of households in contact with multiple sellers will be small, and so the equilibrium distribution of sellers' prices will be bunched up close to the monopoly price with relatively little price dispersion. For goods that account for a large share of household expenditures, the fraction of households in contact with multiple sellers will be large and so the equilibrium distribution of sellers' prices will be bunched up close to the marginal cost, also with relatively little price dispersion. Price dispersion will be highest for goods that account for an intermediate share of household expenditures.

4. PRICE DISPERSION ACROSS HOUSEHOLDS

In Subsection 3.2, we showed that there is substantial variation in the prices paid for identical goods in a given market and quarter. Price dispersion may or may not imply that different households systematically pay different prices for an identical bundle of goods. For instance, given the variation in prices we document in Table 2, a household that buys each of its goods at the 90th percentile of the price distribution would spend 72% more than a household that buys each of its goods at the 10th percentile of the price distribution. In contrast, if households were to draw the price of each good independently from the distribution of transaction prices, then all households would pay approximately the same price for their consumption bundles.

In this section, we assess the extent to which price dispersion remains after aggregating goods into baskets by measuring dispersion in household price indexes. We then decompose the dispersion in household prices indexes in order to understand what are the characteristics of households that pay lower prices and what are the mechanisms through which these households achieve these lower prices.

4.1. Measuring Price Dispersion across Households. We construct a price index for each household i following the same approach as in Aguiar and Hurst (2007). First, let $X_{i,t}$ denote the expenditure of household i in period t and let $\bar{X}_{i,t}$ denote the corresponding hypothetical expenditure had the household purchased its basket of goods at the average price paid for each good by all households in market $m(i)$ in period t . Formally, $X_{i,t}$ and $\bar{X}_{i,t}$ are given by

$$(10) \quad X_{i,t} = \sum_{j,k \in i \cap t} P_{j,k} q_{j,k},$$

$$(11) \quad \bar{X}_{i,t} = \sum_{j,k \in i \cap t} \bar{P}_{j,m(i),t} q_{j,k},$$

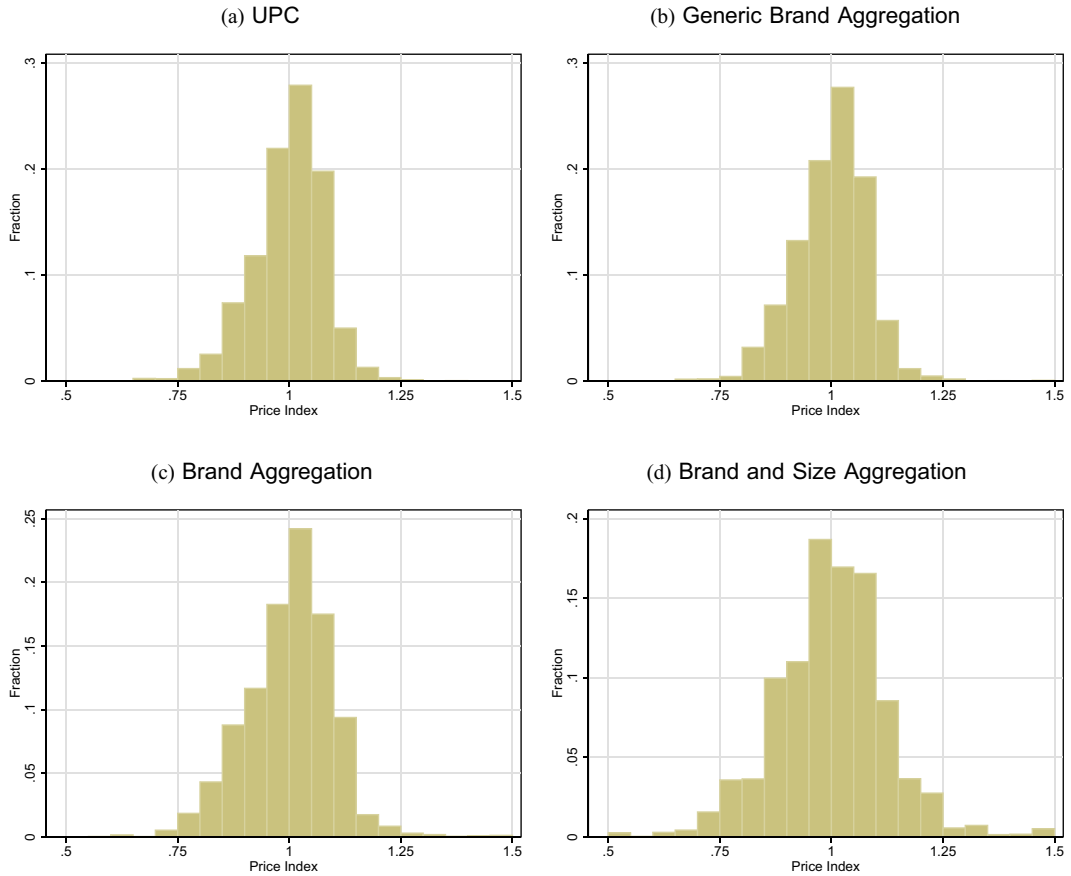
where the notation $k \in i \cap t$ refers to $\{k : t(k) = t, i(k) = i\}$, the set of shopping trips made by household i in quarter t .

The household-specific relative price index, $p_{i,t}$, is defined as the ratio of the actual expenditure of household i in period t to its hypothetical expenditure at market-average prices. That is, $p_{i,t}$ is defined as

$$(12) \quad p_{i,t} = \frac{X_{i,t}}{\bar{X}_{i,t}}.$$

We compute the distribution of price indexes across households separately for each market and quarter. As a concrete example, Figure 7 shows histograms of these price indexes in Minneapolis in 2007:Q1 for our four definitions of a good. These graphs illustrate several features of the distribution of household price indexes. First, for all definitions of a good, there is a significant amount of price index dispersion across households. Second, as we consider broader definitions of a good, the amount of price index dispersion across households increases. Third, there is less dispersion in household price indexes than there is dispersion in transaction prices (which can be seen by comparing Figures 7 and 2).

These features of the distribution of household price indexes in Minneapolis in 2007:Q1 are typical of the distribution of household price indexes in general. The expenditure weighted means of statistics of the household price index distribution (across markets and quarters) are reported in Table 6 for each definition of a good. When goods are defined by their UPC, the average standard deviation of household price indexes is 9%, which is approximately half of the average standard deviation of transaction prices (19%). The average 90–10 percentile ratio of the distribution of household price indexes is 1.22, which is substantially lower than the average



NOTES: Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1.

FIGURE 7

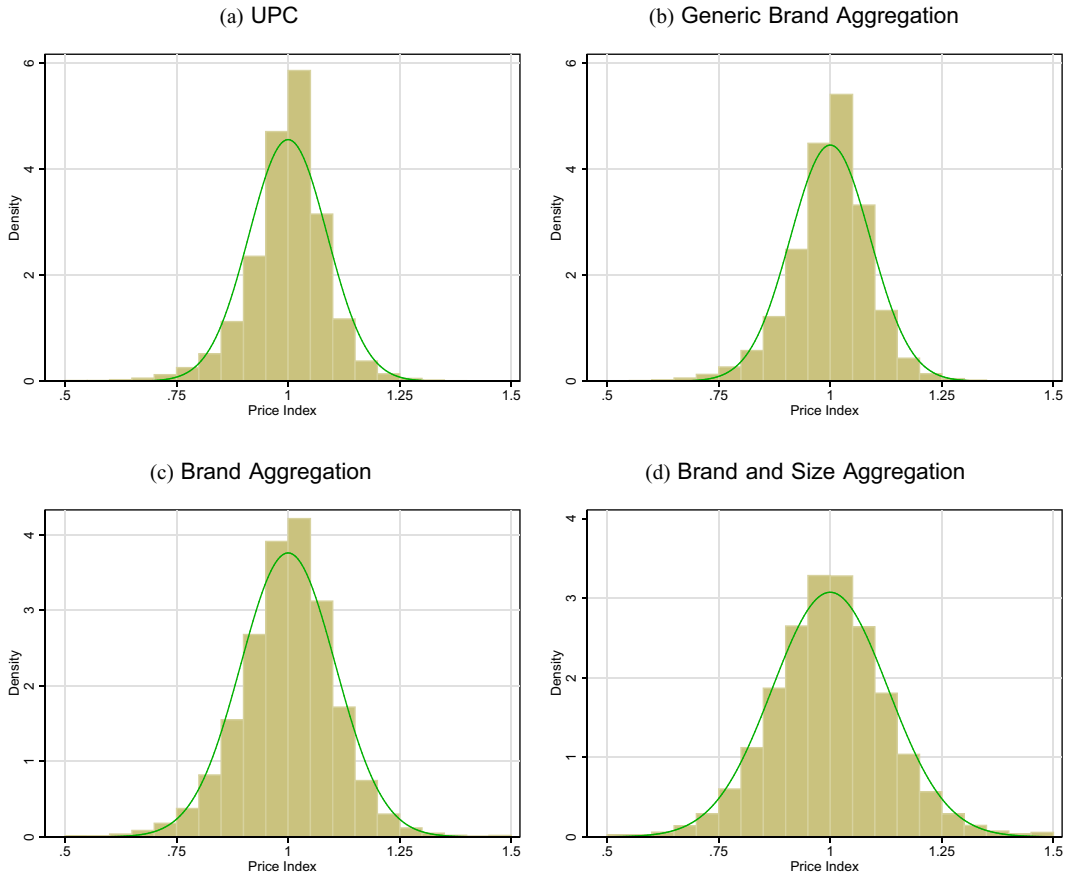
DISTRIBUTION OF PRICE INDEXES IN MINNEAPOLIS, 2007:Q1

TABLE 6
AVERAGE STATISTICS OF PRICE INDEX DISTRIBUTION

	UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Standard deviation	0.09	0.09	0.11	0.13
90–10 ratio	1.22	1.24	1.30	1.38
90–50 ratio	1.09	1.09	1.12	1.16
50–10 ratio	1.12	1.13	1.16	1.19

NOTES: All statistics are expenditure weighted by household and time periods. Percentile ratios are means across different markets and time periods. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1.

90–10 percentile ratio in the distribution of transaction prices (1.72). When goods are defined more broadly, both the average standard deviation and the average 90–10 percentile ratio of the distribution of household price indexes increase, but note that they are always substantially lower than their counterparts for the distribution of transaction prices. Hence, once goods are aggregated into household-level consumption baskets, a large fraction of the variation in prices for individual goods cancels out. Even so, substantial dispersion remains: A household at the 90th percentile of the distribution of price indexes pays 20%–40% more for an identical basket



NOTES: Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1. Overlaid curve is the density of a normal distribution with the same mean and variance.

FIGURE 8

AVERAGE PRICE INDEX DISTRIBUTION

of goods, than a household at the 10th percentile of the distribution, depending on the particular definition of a good.

Figure 8 plots the average distribution of households' price indexes across all markets and quarters, for each definition of a good. The figure reveals that, for all definitions, the 'typical' price index distribution is symmetric, with a unique mode and a distribution that is close in shape to a Normal distribution with the same mean and variance.

4.2. Deconstructing Price Dispersion across Households. Consider two households who purchase the same basket of goods at different prices in the same market and time period. The difference in the prices paid by these two households may originate from three conceptually different sources of price dispersion. First, the two households may have paid different prices because one household purchased its goods at expensive stores, whereas the other household purchased its goods at cheap stores. Second, the two households may have paid different prices because, even though they made all their purchases at equally expensive stores, one household purchased each good in its basket at a store where that particular good was expensive, whereas the other household purchased each good in its basket at a store where that particular good was cheap. Third, the two households may have paid different prices because, even though both households purchased the same goods at exactly the same stores, one household either

purchased more of its goods using coupons or purchased more of its goods on days when those goods were temporarily on sale.

To measure the fraction of price index dispersion that is due to each of these sources, we use the decomposed prices from Subsection 3.3 to construct an analogous decomposition of household price indexes. Let $\omega_{i,j,k,t}$ denote the share of household i 's expenditure in period t that comes from the purchase of good j on shopping trip k , using market average prices,

$$(13) \quad \omega_{i,j,k,t} = \frac{\bar{P}_{j,m(i),t} q_{j,k}}{\bar{X}_{i,t}}.$$

By summing over the purchases of household i in period t using the weights $\omega_{i,j,k,t}$, we can write the household's price index $p_{i,t}$ as

$$(14) \quad p_{i,t} = \sum_{j,k \in i \cap t} \mu_{j,m,t} \omega_{i,j,k,t} + \sum_{j,k \in i \cap t} (\mu_{s,t} - \mu_{j,m,t}) \omega_{i,j,k,t} \\ + \sum_{j,k \in i \cap t} (\mu_{j,s,t} - \mu_{s,t}) \omega_{i,j,k,t} + \sum_{j,k \in i \cap t} (p_{j,k} - \mu_{j,s,t}) \omega_{i,j,k,t},$$

where it is implicit that the market m and store s refer to $m(i)$ and $s(k)$ respectively, i.e., the market in which household i is located and the store corresponding to shopping trip k .

The first term on the right-hand side of (14) is the market average price of good j , averaged over all the goods purchased by household i in quarter t . This term is always equal to 1, since prices have been normalized so that the mean price of any good j in any market m and quarter t is 1. The second term on the right-hand side of (14) is the weighted average of the store components of the transaction prices paid by household i in quarter t . This is the store component of the price index. The third term on the right-hand side of (14) is the weighted average of the store-specific good components of the prices paid by household i in quarter t . This is the store-specific good component of the price index. The fourth term on the right-hand side of (14) is the weighted average of the transaction components of the prices paid by household i in quarter t . This is the transaction component of the price index.

Using the identity in (14), we can decompose the variance of household price indexes into the variances of the store, store-specific good, and transaction components and three covariance terms. As with the decomposition of transaction prices in Table 3, it is important to account for the scarcity of data for some store/good combinations. We use the approach of restricting attention to stores and goods with a minimum of five transactions, since the alternative approach, which imposes structure on the variance of the transaction component, does not allow estimation of the store-specific good component for each store and good and so cannot be used to decompose household price indexes.

The results of this decomposition are reported in Table 7 for the four definitions of a good. For the Generic Brand Aggregation definition (column (2)) the store component accounts for approximately 39% of the variance of price indexes, the store-specific good component accounts for approximately 53%, and the transaction component accounts for approximately 16%. The covariance terms, which in total are negative, constitute -8% of the overall variance of price indexes. The decompositions for the other three definitions of a good (columns (1), (3), and (4)) are similar.

The decompositions in Tables 7 and 3 reveal that dispersion in household price indexes arises from very different sources compared with dispersion in transaction price. First, whereas the store component contributes only around 10% to the variance of transaction prices, it contributes around 40% to the variance of household price indexes. Hence, even though only a small fraction of the dispersion in the transaction prices for the same good is due to differences in the expensiveness of the stores where the transactions took place, nearly half of the dispersion in the prices paid by different households for the same consumption basket is due to differences

TABLE 7
DECOMPOSITION OF VARIANCE OF HOUSEHOLD PRICE INDEXES

	UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Store-specific component	42%	39%	36%	42%
Store-good component	58%	53%	47%	45%
Transaction component	19%	16%	16%	19%
Covar: store, store-good	-24%	-13%	-7%	-16%
Covar: store, transaction	0%	1%	1%	2%
Covar: store-good, transaction	4%	5%	6%	8%

NOTES: All statistics are expenditure weighted across markets and quarters. Sample includes only transactions at stores with unique identifiers for goods/markets/quarters with at least 25 transactions and coefficient of variation less than 1, restricted to store/good combinations with at least five transactions. See text for details of decomposition. May not add to 100% due to rounding.

in the expensiveness of the stores where the households shopped. Second, Table 3 shows that the transaction component is at least as important as the store-good component for explaining the variance of transaction prices. In contrast, Table 7 shows that the store-good component is far more important than the transaction component for explaining the variance of household price indexes. These findings suggest that while there is a great deal of variation in the price at which the same store sells the same good on different days (or in different transactions on the same day), there is not a great deal of variation in households' abilities to systematically take advantage of these differences by properly timing their purchases. In contrast, there does seem to be significant variation in households' abilities to systematically take advantage of persistent price differences for the same good at different stores by purchasing each good at the store where that particular good is, on average, cheaper.

The differences in the extent and sources of dispersion in transaction prices and household price indexes can be rationalized by a simple shopping model. Consider an environment in which there are a large number of stores, each of which sells N goods, and a large number of households who each purchase a basket consisting of all N goods. The households are of two types: regular households who each visit only one store where they purchase all their goods and bargain hunting households who visit two or three stores and purchase each good from the store where it is, on average, cheaper. We assume that both types of households are similar in their abilities to exploit temporary price reductions by visiting their chosen stores more frequently. This shopping model is consistent with the observation that in our data each household concentrates its shopping in a very small number of stores (on average 2.3 per quarter).

This shopping model rationalizes the three key differences between the extent and sources of dispersion of transaction prices and household price indexes: (i) there is less variation in household price indexes than in transaction prices, (ii) the store component accounts for a larger fraction of the variance of household price indexes than of transaction prices, and (iii) the transaction component accounts for a smaller fraction of the variance of household price indexes than of transaction prices. In the shopping model, differences in the store component of transaction prices do not average out at the household level, since each household buys the same basket of goods from a small subset of stores. Differences in the store-good component of prices will average out among households of the same type, but not across the two types of households, since bargain hunting households can purchase each good at the store where that good is on average cheaper, whereas regular households have to buy each good from the same store. Differences in the transaction component of prices will tend to average out among households of the same type and will be very small across household types due to our assumption that households are all similar in their abilities to take advantage of temporary price reductions.

TABLE 8
EFFECT OF SHOPPING BEHAVIOR ON HOUSEHOLD PRICE INDEXES

	(1)	(2)	(3)	(4)
No. shopping trips ($\times 10^2$)	-0.144** (0.005)			-0.035** (0.004)
No. stores visited ($\times 10^2$)		-1.063** (0.027)		-0.556** (0.022)
Fraction of transactions involving coupons			-0.324** (0.003)	-0.317** (0.003)
Observations	880,104	880,104	880,104	880,104
Households	78,758	78,758	78,758	78,758
R^2	0.015	0.021	0.181	0.187

NOTES: All regressions control for time dummies, market dummies, cubic polynomial in the average age of household heads, quadratic polynomial in log household size, and education dummies. Controlling also for expenditure at average prices yields similar results. Dependent variable is household price index using Generic Brand Aggregation definition. Standard errors in parentheses. ** denotes significance at 1%.

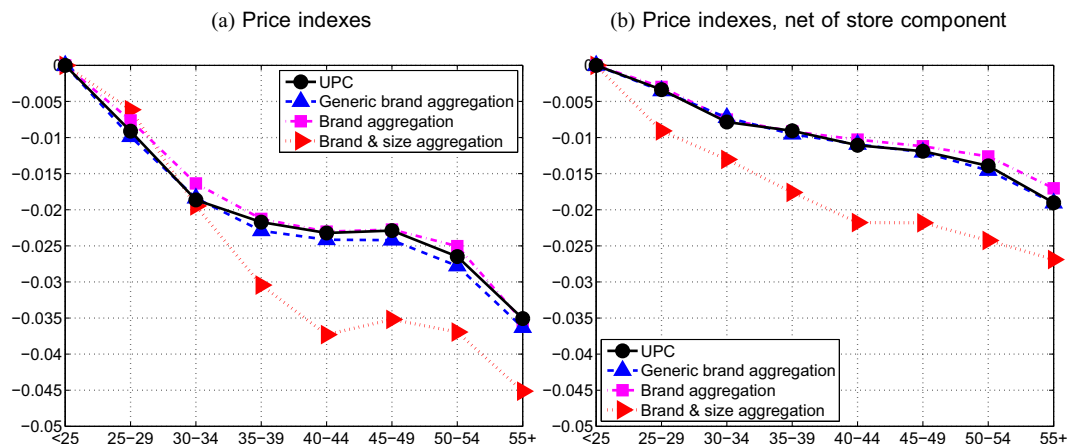
To further understand the features of the household shopping process, we regress households' price indexes on various aspects of their shopping behavior. The results are reported in Table 8 for the Generic Brand Aggregation definition of a good.¹⁸ Columns (1) and (2) show that price indexes are lower both for households that go on more shopping trips and for households who visit a greater number of stores. Column (3) shows that households who use coupons more frequently pay lower prices. Column (4) shows that these findings hold conditionally. That is, households who take more shopping trips pay a lower price for their consumption basket even after controlling for the number of stores they visit and for their coupon usage. Similarly, households who visit more stores or use coupons more frequently pay a lower price for their consumption basket even after controlling for the other two aspects of shopping behavior. The signs in the regressions are consistent with the shopping model described above. Moreover, column (4) implies that the effect on a household's price index of visiting one additional store (0.6%) is an order of magnitude larger than the effect of visiting the same store one more time (0.04%), which lends support to the hypothesis that households find it hard to use multiple shopping trips to exploit temporary sales.¹⁹

4.3. Who Pays Low Prices? The substantial dispersion in prices paid by different households for the same goods begs the question of which types of households enjoy lower prices. The findings in Subsection 4.2 suggest that there are potentially large gains either from being better informed about the location of cheap goods, from being able to take advantage of temporary sales, or from shopping with coupons. Thus it is likely that households whose value of time is lower are better positioned to purchase at lower prices. In this section, we explore this hypothesis by examining the effect of two observable proxies for the value of time: age and employment.

4.3.1. Age and prices. We measure the age of a household as the average age of its heads and plot the age profile of average price indexes in Figure 9(a). We group households into five-year age bins and normalize the average price index of households under 25 years old to zero. The figure illustrates clearly that average household price indexes decline with age. This figure is consistent with the findings in Aguiar and Hurst (2007), who use a small subset of the KNCP data to show that older households pay lower prices than younger households. Figure 9(a) also

¹⁸ The regressions in Table 8 do not control for expenditure at average prices; however, controlling for expenditure does not change any of the conclusions. The results are reported for the Generic Brand Aggregation definition only, but these findings are not sensitive to the particular definition of a good.

¹⁹ The fact that households do not differ much in their ability to take advantage of temporary sales may suggest that either sales are not very effective ways to intertemporally discriminate different types of buyers or that some sales are responses to changes in sellers' inventories (as suggested by Aguirregabiria, 1999). We are grateful to an anonymous referee for pointing this out.



NOTES: Age refers to average age of household head(s).

FIGURE 9

AVERAGE PRICE INDEXES BY AGE

shows that the age profile of household price indexes is almost identical whether we use the UPC, Generic Brand Aggregation, or Brand Aggregation definitions of a good. In contrast, when goods are defined using the Brand and Size Aggregation definition, the life cycle decline in the average household price index is steeper. Indeed, for the three narrower definitions of a good, the average price index for households over 55 years old is 3.5% lower than the average price index for households under 25 years old. For the broadest definition of a good, the average price index for households over 55 years old is 4.5% lower than the average price index for households under 25 years old.

One interpretation of the decline in price indexes over the life cycle is that older households have a lower opportunity cost of time and hence either spend more time searching for low prices, wait for temporary sales, or become better informed about the locations of cheap sellers of particular goods. An alternative interpretation is that older households care less about amenities when shopping and hence shop at stores that offer fewer services and charge lower prices. In order to distinguish between these two interpretations, Figure 9(b) plots the life cycle profile of household price indexes net of the store component, which we argued in Subsection 3.4 is an upper bound on the component of prices that reflects the consumption of amenities. The figure shows that, for all definitions of a good, the average household price index net of the store component declines over the life cycle. However, the decline in the average price index net of the store component is approximately 50% smaller than the decline in the average price index when the store component is included. This implies that, at most, half of the difference in prices paid between old and young households could be due to differences in the consumption of amenities.

4.3.2. *Employment and prices.* To examine the differences in prices paid across households with different employment compositions, we estimate the following regression:

$$(15) \quad \log p_{i,t} = \alpha' e_{i,t} + \beta' x_{i,t} + \lambda_{m(i)} + \kappa_t + \epsilon_{i,t}.$$

We estimate (15) separately for households with one head and households with two heads. The dependent variable, $\log p_{i,t}$, is the logarithm of the price index for household i in quarter t . Our main explanatory variable is $e_{i,t}$, which is a measure of the household's employment composition. For households with one head, $e_{i,t}$ is a dummy variable for whether the head is not employed. For households with two heads, $e_{i,t}$ is a vector of dummy variables, one for

TABLE 9
EFFECT OF EMPLOYMENT ON HOUSEHOLD PRICE INDEXES

	Definition of Good			
	UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Panel A: Households with one head				
Not employed	-0.008** (0.002)	-0.009** (0.002)	-0.019** (0.003)	-0.026** (0.004)
Observations	132,274	137,079	141,505	145,264
Households	15,023	15,201	15,384	15,523
R ²	0.009	0.011	0.017	0.023
Panel B: Households with two heads				
Both heads not employed	-0.009* (0.004)	-0.013** (0.005)	-0.028** (0.005)	-0.046** (0.006)
One head not employed	-0.007** (0.001)	-0.007** (0.001)	-0.011** (0.002)	-0.015** (0.002)
Observations	292,634	299,580	306,505	312,950
Households	33,758	34,072	34,436	34,735
R ²	0.006	0.007	0.013	0.012

NOTES: All regressions control for time dummies, market dummies, cubic polynomial in the average age of household heads, quadratic polynomial in log household size, and education dummies. Dependent variable is household price index. Standard errors in parentheses. * denotes significance at 5%, ** denotes significance at 1%.

whether exactly one of the heads is not employed and one for whether both of the heads are not employed. The remaining explanatory variables are a set of household-level characteristics $x_{i,t}$ that include age, household size, and education dummies as well as a full set of time κ_t and market $\lambda_{m(t)}$ dummy variables. In order to focus on the difference between employment and nonemployment, instead of between employment and retirement, we restrict attention to households whose heads' average age is between 25 and 55.²⁰

We report estimates of Equation (15) for households with one head in Panel A of Table 9. The estimates indicate that one-head households who are not employed pay between 0.8% and 2.6% less for the same basket of goods than one-head households who are employed after controlling for differences in age, household size, and education. The estimates for households with two heads are reported in Panel B of Table 9. Households with two nonemployed heads pay between 0.9% and 4.6% less than households with two employed heads. Households with one nonemployed head pay between 0.7% and 1.5% less than households with two employed heads. In all cases, the effects are smallest when we define goods by their UPC and become larger as the definition of a good is broadened. All of the estimates are strongly statistically significant, and suggest that the effect of employment on price indexes is roughly the same magnitude as the fall in price indexes over the life cycle.²¹

²⁰ The information on employment refers to the employment status at the time of the survey. The survey does not distinguish between unemployed individuals and nonparticipants, but does explicitly identify students who work less than 30 hours per week. Age and household size are included as explanatory variables in order to control for differences in purchasing needs across households. We have experimented with controlling for education as a proxy for permanent income differences between the employed and nonemployed and controlling for household expenditure at average prices as an additional proxy for needs.

²¹ Recall that the demographic information in the KNCP is updated only once a year in the last quarter, usually in December. Since employment transitions frequently occur within a year or even a quarter, our measure of employment is possibly mismeasured in the later months of the year. To assess whether this potential measurement error is important, we have estimated equivalent models on a subsample that includes only transactions in the first quarter of a year. Because of the timing of the demographic questions, reported employment status is much more likely to be accurate in the first few months of each calendar year than later in the year. We find that our results are barely affected by this restriction, suggesting that this type of measurement error is unlikely to be influencing our findings in an important way.

TABLE 10
EFFECT OF EMPLOYMENT ON HOUSEHOLD PRICE INDEXES NET OF STORE COMPONENT

	Definition of Good			
	UPC (1)	Generic Aggregation (2)	Brand Aggregation (3)	Brand & Size Aggregation (4)
Panel A: Households with one head				
Not employed	-0.006** (0.002)	-0.006** (0.002)	-0.011** (0.002)	-0.012** (0.003)
Observations	132,274	137,079	141,505	145,264
Households	15,023	15,201	15,384	15,523
R ²	0.010	0.009	0.008	0.009
Panel B: Households with two heads				
Both heads not employed	-0.007** (0.003)	-0.011** (0.003)	-0.022** (0.004)	-0.035** (0.006)
One head not employed	-0.004** (0.001)	-0.004** (0.001)	-0.007** (0.001)	-0.009** (0.002)
Observations	292,634	299,580	306,505	312,950
Households	33,758	34,072	34,436	34,735
R ²	0.007	0.008	0.010	0.007

NOTES: All regressions control for time dummies, market dummies, cubic polynomial in the average age of household heads, quadratic polynomial in log household size, and education dummies. Dependent variable is household price index net of the store component. Standard errors in parentheses. ** denotes significance at 1%.

As with older households, one interpretation of the findings in Table 9 is that nonemployed households have a lower opportunity cost of time and hence spend more time searching for low prices or waiting for temporary sales. An alternative interpretation is that nonemployed households care less about amenities when shopping and hence shop at stores that offer fewer amenities and charge lower prices. In order to distinguish between these two interpretations, we reestimate Equation (15) but replace the dependent variable with the logarithm of the household's price index net of the store component. The estimates, which are reported in Table 10, show that nonemployment has a large and statistically significant effect on the household's price index net of the store component, but that the effects are approximately 30% smaller than the corresponding effects of nonemployment on the households' price index in Table 9. These findings suggest that at most one-third of the difference in prices paid between employed and nonemployed households is due to differences in the consumption of amenities, and at least two-thirds is due to differences in the ability or willingness of households to search for lower prices.

5. CONCLUSION

Our goal in this article has been to use household-level scanner data to undertake a systematic study of the morphology of price distributions, with the aim of providing new insights into the theoretical origins of price dispersion. We found that the distribution of prices for individual goods is leptokurtic with a standard deviation between 19% and 36%, whereas the distribution of prices for bundles of goods is Normal with a standard deviation between 9% and 14%. For both types of prices, dispersion is a robust feature of the data.

We then examined both data and theories through the lens of a variance decomposition that splits prices into store components, store-specific good components, and transaction components. Our decomposition casts doubt on the quantitative importance of differences in amenities and marginal costs as sources of price dispersion, since both imply that the store component should account for observed price dispersion, yet we found that the store component accounts for only 10% of the variance of prices. In contrast, our decomposition

supports intertemporal price discrimination and search frictions as quantitatively important theories of price dispersion, since these theories imply that the store-specific good and transaction components should account for observed price dispersion.

We also examined the characteristics of households who pay lower prices for identical bundles of goods. We found that older households and households with fewer employed members pay lower prices and do so mostly by visiting a greater number of stores on each shopping trip instead of shopping more frequently. These findings suggest heterogeneity in households' costs of time are important for understanding dispersion in household price indexes; this is further evidence in support of the quantitative importance of search frictions.

We see at least two important directions for future theoretical work on pricing, motivated by our empirical findings. First, there are currently no theories that encompass all of the features of price distributions that we document. For example, the striking discrepancy between the decomposition of individual transaction prices and dispersion in household price indexes highlights a need for new theories of shopping behavior and pricing of bundles. Second, our findings indicate that activities that can be described as bargain hunting (e.g., visiting multiple stores in search of the cheapest price for a particular good, waiting for sales, or seeking out coupons) are crucial for understanding the data on price dispersion. Bargain hunting requires time. We have shown that older households and working-age households with fewer employed members—both natural proxies for the value of a household's time—do, in fact, pay lower prices and that they achieve this through bargain hunting activities. A natural next step is to begin understanding the macroeconomic implications of bargain hunting—for example, by integrating search theories of price dispersion into a general equilibrium framework. In Kaplan and Menzio (2014), we develop a theoretical model in this direction.

APPENDIX

A.1. Additional Tables and Figures. This appendix contains additional tables and figures that are referred to in the main text, but which we have excluded from the main body of the article due to space limitations. Table A.1 lists the 54 scantrack markets covered by the KNCP data, together with the fraction of observed expenditure in 2007:Q1 that comes from each market. Table A.2 lists the departments covered by the KNCP data, the number of product groups, and product modules in each department together with the number of different UPCs purchased from each department and the fraction of expenditure coming from each department in 2007:Q1. Table A.3 lists the types of retailers covered by the KNCP data together with the fraction of expenditure coming from each category in 2007:Q1.

A.2. Details of Price Decompositions. In this appendix, we provide further details of the price decompositions and their implementation. Unless otherwise specified, the notation in this appendix suppresses conditioning on the market m and quarter t corresponding to the transaction k . It is implicit that all formulas apply separately to each market and quarter.

A.2.1. Notation. We start by defining the notation used in the appendix, together with some useful identities. Let $P_{j,k}$ be the actual price for good j on shopping trip k , with corresponding quantity $q_{j,k}$. Let $\omega_{j,k}$ denote the quantity weight for good j of shopping trip k , defined as

$$(A.1) \quad \omega_{j,k} = \frac{q_{j,k}}{\sum_k q_{j,k}}.$$

TABLE A.1
FRACTION OF EXPENDITURE FROM EACH SCANTRACK MARKET

Scantrack Market	Transactions		Scantrack Market	Transactions	
	All	With Store ID		All	With Store ID
Albany	0.9%	1.0%	Milwaukee	1.1%	1.3%
Atlanta	2.5%	2.4%	Minneapolis	1.9%	1.7%
Baltimore	1.2%	1.4%	Nashville	1.2%	0.7%
Birmingham	1.3%	0.6%	New Orleans-Mobile	1.6%	0.9%
Boston	3.5%	4.4%	NYC (ex-urban)	1.2%	1.6%
Buffalo-Rochester	1.1%	1.2%	NYC (suburban)	3.6%	4.8%
Charlotte	1.1%	1.2%	NYC (urban)	3.1%	2.6%
Chicago	3.6%	4.0%	Oklahoma City-Tulsa	1.3%	0.4%
Cincinnati	1.5%	1.8%	Omaha	0.7%	0.4%
Cleveland	2.1%	1.8%	Orlando	1.7%	1.8%
Columbus	1.1%	1.3%	Philadelphia	3.8%	4.3%
Dallas	2.7%	2.3%	Phoenix	2.6%	2.8%
Denver	1.9%	1.9%	Pittsburgh	1.8%	1.4%
Des Moines	0.5%	0.4%	Portland	1.7%	1.5%
Detroit	2.7%	3.4%	Raleigh-Durham	1.9%	1.9%
Grand Rapids	1.2%	1.5%	Richmond	1.6%	1.5%
Hartford-New Haven	1.6%	1.8%	Sacramento	1.6%	1.5%
Houston	2.1%	1.9%	Salt Lake City-Boise	1.5%	1.0%
Indianapolis	1.5%	1.5%	San Antonio	1.5%	1.5%
Jacksonville	0.9%	0.8%	San Diego	1.3%	1.4%
Kansas City	1.2%	0.6%	San Francisco	2.9%	3.3%
Las Vegas	0.9%	0.9%	Seattle	2.1%	2.0%
Little Rock	0.8%	0.4%	St. Louis	1.4%	1.3%
Los Angeles	6.4%	7.3%	Syracuse	0.9%	0.8%
Louisville	1.2%	1.1%	Tampa	2.4%	2.7%
Memphis	1.0%	0.5%	Washington, DC	3.1%	3.5%
Miami	2.7%	3.6%	West Texas	1.0%	0.3%

NOTES: Expenditure data refer to 2007:Q1. Expenditure is weighted by household projection factors. Statistics in columns labeled “With Store ID” condition on transactions at stores that have unique store identifiers.

TABLE A.2
DEPARTMENTS, PRODUCT GROUPS, AND PRODUCT MODULES

Department	Product Groups	Product Modules	UPCs	Transactions				With Store ID in Minneapolis	
				UPCs	Frac	UPCs	Frac	UPCs	Frac
Health and beauty	18	173	242,564	68,781	11%	43,996	11%	6,024	12%
Dry grocery	38	416	499,943	187,962	39%	129,324	42%	16,305	39%
Frozen foods	10	85	83,209	32,303	9%	23,773	11%	3,467	12%
Dairy	10	45	63,782	30,202	8%	21,747	10%	2,301	10%
Deli	1	16	23,667	6730	2%	5221	3%	509	2%
Packaged meat	2	13	23,981	9501	3%	6881	3%	854	4%
Fresh produce	1	21	18,378	5993	3%	4539	3%	504	3%
Non-food grocery	12	136	172,308	49,917	14%	27,360	11%	3,624	11%
Alcohol	3	30	40,303	12,029	4%	5791	3%	108	0%
General merch.	18	147	256,283	58,063	7%	24,631	4%	2,408	6%
Total	113	1082	1,424,418	461,481		293,263		36,104	

NOTES: First three columns report number of product groups, product modules, and UPCs in entire database. Remaining columns report number of UPCs transacted and fraction of expenditure in 2007:Q1. Expenditure data are weighted by household projection factors. Statistics in columns labeled “With Store ID” condition on transactions at stores that have unique store identifiers. Statistics in columns labeled “With Store ID in Minneapolis” condition on transactions at stores in Minneapolis that have unique store identifiers.

TABLE A.3
FRACTION OF EXPENDITURE BY TYPE OF RETAILER

Retailer Type	Transactions		
	All	With Store ID	With Store ID in Minneapolis
Grocery	49%	82%	62%
Discount Store	22%	9%	32%
Warehouse Club	10%		
Drug Store	5%	9%	6%
Dollar Store	2%		
Online Shopping	1%		
Liquor Store	1%		
All Other Stores	10%		

NOTES: Expenditure data refer to 2007:Q1. Expenditure is weighted by household projection factors. Statistics in columns labeled “With Store ID” condition on transactions at stores that have unique store identifiers. Statistics in columns labeled “With Store ID in Minneapolis” condition on transactions at stores in Minneapolis that have unique store identifiers.

Let $p_{j,k}$ denote the normalized price for good j on shopping trip k , defined as

$$(A.2) \quad p_{j,k} = \frac{P_{j,k}}{\sum_k P_{j,k} \omega_{j,k}}.$$

Note that the quantity-weighted average of the normalized prices for good j is 1,

$$(A.3) \quad \sum_k p_{j,k} \omega_{j,k} = 1.$$

Let $\omega_{j,s}$ denote the quantity weight for good j of store s , i.e., the relative contribution of store s to the total purchases of good j :

$$(A.4) \quad \omega_{j,s} = \frac{\sum_{k \in s} q_{j,k}}{\sum_k q_{j,k}}.$$

Note the following relationship between $\omega_{j,k}$ and $\omega_{j,s}$:

$$(A.5) \quad \omega_{j,s} = \sum_{k \in s} \omega_{j,k}.$$

Let $R_{j,s}$ denote the revenue from good j at store s , let R_s denote the total revenue at store s , and let $\rho_{j,s}$ denote the revenue weight for good j at store s , i.e., the relative contribution of good j to the total revenues at store s :

$$(A.6) \quad R_{j,s} = \sum_{k \in s} P_{j,k} q_{j,k},$$

$$(A.7) \quad R_s = \sum_j R_{j,s},$$

$$(A.8) \quad \rho_{j,s} = \frac{R_{j,s}}{R_s}.$$

Note that for every store s the sum of the revenue weights across goods is 1:

$$(A.9) \quad \sum_j \rho_{j,s} = 1.$$

A.2.2. Decomposition of prices. In Subsection 3.3, we decompose $p_{j,k}$ using the identity

$$(A.10) \quad p_{j,k} = \mu_j + (\mu_s - \mu_j) + (\mu_{j,s} - \mu_s) + (p_{j,k} - \mu_{j,s}).$$

The terms on the right-hand side of (A.10) are defined as

$$(A.11) \quad \mu_j = \sum_k p_{j,k} \omega_{j,k},$$

$$(A.12) \quad \mu_{j,s} = \sum_{k \in S} p_{j,k} \frac{\omega_{j,k}}{\omega_{j,s}},$$

$$(A.13) \quad \mu_s = \sum_j \mu_{j,s} \rho_{j,s}.$$

We use the following terminology to describe the components in (A.10):

- $(\mu_s - \mu_j)$ is the *store component* of $p_{j,k}$.
- $(\mu_{j,s} - p_{j,k}) = (\mu_{j,s} - \mu_s) + (p_{j,k} - \mu_{j,s})$ is the *within-store component* of $p_{j,k}$.
- $(\mu_{j,s} - \mu_s)$ is the *store-specific good component* of $p_{j,k}$.
- $(p_{j,k} - \mu_{j,s})$ is the *transaction component* of $p_{j,k}$.

The following four identities are useful in the algebra that follows:

$$(A.14) \quad \mu_j = 1,$$

$$(A.15) \quad \mu_{j,s} \omega_{j,s} = \sum_{k \in S} p_{j,k} \omega_{j,k},$$

$$(A.16) \quad \sum_s \mu_{j,s} \omega_{j,s} = 1,$$

$$(A.17) \quad \sum_{k \in S} \mu_{j,s} \omega_{j,k} = \mu_{j,s} \omega_{j,s}.$$

(A.14) follows from (A.3) and (A.11), (A.15) follows from rearranging (A.12), (A.16) follows from summing (A.15) over s and applying (A.3), and (A.17) follows from (A.5).

We now derive the quantity-weighted means, which we denote by $m(\bullet)$, of each component on the right-hand side of (A.10), for each good j . Note from (A.3) that the quantity-weighted mean of $p_{j,k}$ is 1.

- Store component:

$$\begin{aligned}
 (A.18) \quad m(\mu_s - \mu_j) &= \sum_k (\mu_s - \mu_j) \omega_{j,k} \\
 &= \sum_k \mu_s \omega_{j,k} - 1 \\
 &= \sum_s \mu_s \sum_{k \in s} \omega_{j,k} - 1 \\
 &= \sum_s \mu_s \omega_{j,s} - 1,
 \end{aligned}$$

where the second line uses (A.14) and the fourth line uses (A.5).

- Store-specific good component:

$$\begin{aligned}
 (A.19) \quad m(\mu_{j,s} - \mu_s) &= \sum_k (\mu_{j,s} - \mu_s) \omega_{j,k} \\
 &= \sum_s \sum_{k \in s} \mu_{j,s} \omega_{j,k} - \sum_s \sum_{k \in s} \mu_s \omega_{j,k} \\
 &= \sum_s \mu_{j,s} \omega_{j,s} - \sum_s \mu_s \omega_{j,s} \\
 &= 1 - \sum_s \mu_s \omega_{j,s},
 \end{aligned}$$

where the third line uses (A.17) and the fourth line uses (A.16).

- Transaction component:

$$\begin{aligned}
 (A.20) \quad m(p_{j,k} - \mu_{j,s}) &= \sum_k (p_{j,k} - \mu_{j,s}) \omega_{j,k} \\
 &= 1 - \sum_k \mu_{j,s} \omega_{j,k} \\
 &= 1 - \sum_s \mu_{j,s} \omega_{j,s} \\
 &= 0,
 \end{aligned}$$

where the third line uses (A.5) and the fourth line uses (A.16).

Next we take the quantity-weighted variance of both sides of (A.10) for each good j . For each good j , the quantity-weighted variance of the normalized price $p_{j,k}$ is

$$\begin{aligned}
 \text{var}(p_{j,k}) &= \sum_k p_{j,k}^2 \omega_{j,k} - \left(\sum_k p_{j,k} \omega_{j,k} \right)^2 \\
 &= \sum_k p_{j,k}^2 \omega_{j,k} - 1,
 \end{aligned}$$

where the second line uses (A.3).

The quantity-weighted variances of each of the components on the right-hand side of (A.10) are as follows:

- Store component:

$$(A.21) \quad \begin{aligned} \text{var}(\mu_s - \mu_j) &= \sum_k (\mu_s - \mu_j)^2 \omega_{j,k} - m(\mu_s - \mu_j)^2 \\ &= \sum_k (\mu_s - 1)^2 \omega_{j,k} - \left(\sum_s \mu_s \omega_{j,s} - 1 \right)^2, \end{aligned}$$

where the second line uses (A.14) and (A.18).

- Store-good component:

$$(A.22) \quad \begin{aligned} \text{var}(\mu_{j,s} - \mu_s) &= \sum_k (\mu_{j,s} - \mu_s)^2 \omega_{j,k} - m(\mu_{j,s} - \mu_s)^2 \\ &= \sum_k (\mu_{j,s} - \mu_s)^2 \omega_{j,k} - \left(1 - \sum_s \mu_s \omega_{j,s} \right)^2, \end{aligned}$$

where the second line uses (A.19).

- Transaction component:

$$(A.23) \quad \begin{aligned} \text{var}(p_{j,k} - \mu_{j,s}) &= \sum_k (p_{j,k} - \mu_{j,s})^2 \omega_{j,k} - m(p_{j,k} - \mu_{j,s})^2 \\ &= \sum_k (p_{j,k} - \mu_{j,s})^2 \omega_{j,k}, \end{aligned}$$

where the second line uses (A.20).

The quantity-weighted covariances of each of the components on the right-hand side of (A.10) are as follows:

- Store-specific good component, transaction component:

$$(A.24) \quad \begin{aligned} \text{cov}(\mu_{j,s} - \mu_s, p_{j,k} - \mu_{j,s}) &= \sum_k (\mu_{j,s} - \mu_s)(p_{j,k} - \mu_{j,s}) \omega_{j,k} - m(\mu_{j,s} - \mu_s) m(p_{j,k} - \mu_{j,s}) \\ &= \sum_k (\mu_{j,s} - \mu_s)(p_{j,k} - \mu_{j,s}) \omega_{j,k} \\ &= \sum_s (\mu_{j,s} - \mu_s) \sum_{k \in S} p_{j,k} \omega_{j,k} - \sum_s (\mu_{j,s} - \mu_s) \mu_{j,s} \omega_{j,s} \\ &= \sum_s (\mu_{j,s} - \mu_s) \mu_{j,s} \omega_{j,s} - \sum_s (\mu_{j,s} - \mu_s) \mu_{j,s} \omega_{j,s} \\ &= 0, \end{aligned}$$

where the second line uses (A.20), the third line uses (A.5), and the fourth line uses (A.15).

- Store component, transaction component:

$$(A.25) \quad \begin{aligned} \text{cov}(\mu_s - 1, p_{j,k} - \mu_{j,s}) &= \sum_k (\mu_s - 1)(p_{j,k} - \mu_{j,s}) \omega_{j,k} - m(\mu_s - 1) m(p_{j,k} - \mu_{j,s}) \\ &= \sum_k (\mu_s - 1)(p_{j,k} - \mu_{j,s}) \omega_{j,k} \\ &= \sum_s (\mu_s - 1) \sum_{k \in S} p_{j,k} \omega_{j,k} - \sum_s (\mu_s - 1) \mu_{j,s} \sum_{k \in S} \omega_{j,k} \end{aligned}$$

$$\begin{aligned}
&= \sum_s (\mu_s - 1) \mu_{j,s} \omega_{j,s} - \sum_s (\mu_s - 1) \mu_{j,s} \omega_{j,s} \\
&= 0,
\end{aligned}$$

where the second line uses (A.20) and the fourth line uses (A.15) and (A.5).

- Store component, store-specific good component:

$$\begin{aligned}
\text{cov}(\mu_s - 1, \mu_{j,s} - \mu_s) &= \sum_k (\mu_s - 1) (\mu_{j,s} - \mu_s) \omega_{j,k} - m (\mu_s - 1) m (\mu_{j,s} - \mu_s) \\
&= \sum_s (\mu_s - 1) (\mu_{j,s} - \mu_s) \sum_{k \in s} \omega_{j,k} - \left(\sum_s \mu_s \omega_{j,s} - 1 \right) \left(1 - \sum_s \mu_s \omega_{j,s} \right) \\
\text{(A.26)} \quad &= \sum_s (\mu_s - 1) (\mu_{j,s} - \mu_s) \omega_{j,s} + \left(\sum_s \mu_s \omega_{j,s} - 1 \right)^2 \\
&= \sum_s (\mu_{j,s} - \mu_s) \mu_s \omega_{j,s} + \sum_s \mu_s \omega_{j,s} - 1 + \left(\sum_s \mu_s \omega_{j,s} - 1 \right)^2 \\
&= \sum_s (\mu_{j,s} - \mu_s) \mu_s \omega_{j,s} - \sum_s \mu_s \omega_{j,s} + \left(\sum_s \mu_s \omega_{j,s} \right)^2,
\end{aligned}$$

where the second line uses (A.18) and (A.19), the third line uses (A.5), and the fourth line uses (A.16). In general (A.26) is not equal to zero.

This establishes the claim in Subsection 3.3 that the covariance between the transaction component and the store component, and the covariance between the store-specific good component and the transaction component, are both zero.

A.2.3. Implementing the price decomposition. The variance decomposition is implemented separately for each good j in each market m and quarter t . We require some additional notation in order to explain how we aggregate across goods, markets, and quarters. Let $X_{j,m,t}$ denote the total market expenditure for good j in market m and quarter t , and let $X_{m,t}$ denote the total market expenditure in market m and quarter t . Let $x_{j,m,t}$ denote the expenditure weight for good j , market m , and quarter t , i.e., the contribution of good j in market m and quarter t to the total expenditure across all goods, markets, and quarters, and let $x_{m,t}$ denote the expenditure weight for market m and quarter t , i.e., the contribution of market m and quarter t to the total expenditure across all goods, markets, and quarters:

$$\text{(A.27)} \quad X_{j,m,t} = \sum_{k \in m \cap t} P_{j,k} q_{j,k},$$

$$\text{(A.28)} \quad X_{m,t} = \sum_j X_{j,m,t},$$

$$\text{(A.29)} \quad x_{j,m,t} = \frac{X_{j,m,t}}{\sum_{m,t} X_{m,t}},$$

$$\text{(A.30)} \quad x_{m,t} = \frac{X_{m,t}}{\sum_{m,t} X_{m,t}}.$$

To implement the decomposition, we take the following steps: First, we compute the variances and covariances above for each good in each market and quarter using the formulas in (A.21)–(A.26). Let $\text{var}_{j,m,t}(\bullet)$ denote any of these variances or covariances for a specific good, market, and time period.

Second, we multiply each variance or covariance estimate by a Bessel correction factor that is based on quantity weights, so that our estimates are unbiased. The correction factor for good j in market m and quarter t is defined as

$$(A.31) \quad B_{j,m,t}^q = \frac{\left(\sum_{k \in m \cap t} q_{jk} \right)^2}{\left(\sum_{k \in m \cap t} q_{jk} \right)^2 - \sum_{k \in m \cap t} q_{jk}^2}.$$

Third, for each market and quarter, we compute a weighted average of the variance and covariance of each component across goods. We use weights $b_{j,m,t}^x$ that correspond to the Bessel correction factors based on expenditure so that our estimates are unbiased estimates of common variances and covariances across goods. The weights are defined as

$$(A.32) \quad B_{j,m,t}^x = \frac{X_{j,m,t}^2}{X_{j,m,t}^2 - \sum_{k \in m \cap t} (P_{j,k} q_{jk})^2},$$

$$(A.33) \quad b_{j,m,t}^x = \frac{B_{j,m,t}^x}{\sum_j B_{j,m,t}^x}.$$

Thus the average variances and covariances across all goods in market m at quarter t are given by

- Overall variance of normalized prices:

$$(A.34) \quad \text{var}_{m,t}(p_{j,k}) = \sum_j B_{j,m,t}^q \text{var}_{j,m,t}(p_{j,k}) b_{j,m,t}^x$$

- Variance of store component:

$$(A.35) \quad \text{var}_{m,t}(\mu_s - \mu_j) = \sum_j B_{j,m,t}^q \text{var}_{j,m,t}(\mu_s - \mu_j) b_{j,m,t}^x$$

- Variance of store-specific good component:

$$(A.36) \quad \text{var}_{m,t}(\mu_{j,s} - \mu_s) = \sum_j B_{j,m,t}^q \text{var}_{j,m,t}(\mu_{j,s} - \mu_s) b_{j,m,t}^x$$

- Variance of transaction component:

$$(A.37) \quad \text{var}_{m,t}(p_{j,k} - \mu_{j,s}) = \sum_j B_{j,m,t}^q \text{var}_{j,m,t}(p_{j,k} - \mu_{j,s}) b_{j,m,t}^x$$

- Covariance: store component, store-specific good component:

$$(A.38) \quad \text{cov}_{m,t}(\mu_s - \mu_j, \mu_{j,s} - \mu_s) = \sum_j B_{j,m,t}^q \text{cov}_{j,m,t}(\mu_s - \mu_j, \mu_{j,s} - \mu_s) b_{j,m,t}^x.$$

Fourth, we compute the fraction of variation due to each source separately for each market and quarter by dividing (A.36)–(A.38) by (A.34). We then take the expenditure-weighted average of these fractions across all markets and time periods. So, for example, the average contribution of the store component is given by

$$(A.39) \quad \sum_{m,t} \frac{\text{var}_{m,t}(\mu_s - \mu_j)}{\text{var}_{m,t}(p_{j,k})} x_{m,t}.$$

The contribution of the other components are defined analogously.

A.2.4. Restriction on the variance of the transaction component. In Subsection 3.3, we decompose prices under the assumption that the variance of the transaction component is constant across certain subgroups of goods and markets. The procedure is as follows: First, we assume that the variance of the transaction component for good j in sold at store s in market m in quarter t is given by

$$(A.40) \quad \text{var}_{j,s,t}(p_{j,k} - \mu_{j,s}) = \alpha_{D(j),t} + \alpha_{m(s),t} + \varepsilon_{j,s,t}$$

where $\text{var}_{j,s,t}(p_{j,k} - \mu_{j,s})$ refers to the variance of the transaction component of normalized prices for good j at store s :

$$(A.41) \quad \text{var}_{j,s,t}(p_{j,k} - \mu_{j,s}) = \sum_{k \in s} (p_{j,k} - \mu_{j,s})^2 \omega_{j,k} - \left(\sum_{k \in s} (p_{j,k} - \mu_{j,s}) \omega_{j,k} \right)^2.$$

The first term on the right-hand side of (A.40) is a department effect, where the department corresponding to good j is denoted by $D(j)$. The second term on the right-hand side of (A.40) is a market effect. The third term is an error term.

We estimate (A.40) using weighted least squares, with weights given by the Bessel factors corresponding to the expenditure on good j at store s :

$$(A.42) \quad B_{j,s,m,t}^x = \frac{R_{j,s,m,t}^2}{R_{j,s,m,t}^2 - \sum_{k \in s \cap m \cap t} (P_{j,k} q_{jk})^2}.$$

These weights account for the fact that $\mu_{j,s}$ are sample means that have been estimated with a potentially different number of transactions for each store/good pair. Since (A.40) implies that the predicted variance of the transaction component for each good is the same for all stores in a given market, the predicted variance of the transaction component for market m in quarter t is given by

$$(A.43) \quad \hat{\text{var}}_{j,m,t}(p_{j,k} - \mu_{j,s}) = \hat{\alpha}_{D(j),t} + \hat{\alpha}_{m(s),t}.$$

We then compute the variance of the store–good component as

$$(A.44) \quad \hat{\text{var}}_{j,m,t}(\mu_{j,s} - \mu_s) = \text{var}_{j,m,t}(p_{j,k} - \mu_s) - \hat{\text{var}}_{j,m,t}(p_{j,k} - \mu_{j,s}).$$

A.2.5. *Retailer decomposition of prices.* In Subsection 3.4, we decompose p_{jk} using the two identities

$$(A.45) \quad \mu_s - \mu_j = (\mu_r - \mu_j) + (\mu_s - \mu_r),$$

$$(A.46) \quad p_{j,k} - \mu_s = (\mu_{j,r} - \mu_r) + (p_{j,k} - \mu_s - \mu_{j,r} + \mu_r).$$

To define the terms on the right-hand side of (A.45) and (A.46), we require some more notation. Let $\omega_{j,r}$ denote the quantity weight for good j of retailer r , i.e., the relative contribution of retailer r to the total purchases of good j , defined as

$$(A.47) \quad \omega_{j,r} = \frac{\sum_{k \in r} q_{j,k}}{\sum_k q_{j,k}}.$$

Let $R_{j,r}$ denote the revenue from good j at store r , let R_r denote the total revenue at retailer r , and let $\rho_{j,r}$ denote the revenue weight for good j at store r , i.e., the relative contribution of good j to the total revenues of retailer r :

$$(A.48) \quad R_{j,r} = \sum_{k \in r} P_{j,k} q_{j,k},$$

$$(A.49) \quad R_r = \sum_j R_{j,r},$$

$$(A.50) \quad \rho_{j,r} = \frac{R_{j,r}}{R_r}.$$

The objects on the right-hand side of (A.45) and (A.46) are defined as

$$(A.51) \quad \mu_{j,r} = \sum_{k \in r} p_{j,k} \frac{\omega_{j,k}}{\omega_{j,r}},$$

$$(A.52) \quad \mu_r = \sum_j \mu_{j,r} \rho_{j,r}.$$

The quantity-weighted variances and covariances of each of the terms in (A.45) and (A.46) are defined analogously to those in (A.21)–(A.26). Note that the covariance between the terms in (A.45) and the terms in (A.46) is equal to the covariance between the store component and the store-specific good component in (A.26), since the other covariance terms were shown above to be zero.

A.2.6. *Household price index decomposition.* Let $X_{i,j,k}$ denote the expenditure of household i on good j on shopping trip k , and let X_i denote the total expenditure of household i on all goods,

$$(A.53) \quad X_{i,j,k} = P_{j,k} q_{j,k},$$

$$(A.54) \quad X_i = \sum_{j,k \in i} X_{i,j,k}.$$

Let $\bar{X}_{i,j,k}$ denote the hypothetical expenditure of household i on good j on shopping trip k evaluated at the average price of good j in market $m(i)$ in the corresponding quarter, and let \bar{X}_i denote the corresponding hypothetical total expenditure at average prices of household i on all goods,

$$(A.55) \quad \bar{X}_{i,j,k} = \bar{P}_j q_{j,k},$$

$$(A.56) \quad \bar{X}_i = \sum_{j,k \in i} \bar{X}_{i,j,k},$$

where the average price \bar{P}_j is defined as

$$(A.57) \quad \bar{P}_j = \sum_k P_{j,k} \omega_{j,k}.$$

Let $\omega_{i,j,k}$ be the expenditure weight for household i of good j on shopping trip k , evaluated at average prices, i.e., the contribution of good j on shopping trip to household i 's total expenditure at average prices:

$$(A.58) \quad \omega_{i,j,k} = \frac{\bar{X}_{i,j,k}}{\bar{X}_i}.$$

In Subsection 4.2, we define the price index for household i , p_i , as

$$(A.59) \quad p_i = \frac{X_i}{\bar{X}_i}.$$

We decompose p_i by writing p_i as

$$(A.60) \quad \begin{aligned} p_i &= \frac{X_i}{\bar{X}_i} \\ &= \frac{\sum_{j,k \in i} P_{j,k} q_{j,k}}{\bar{X}_i} \\ &= \frac{\sum_{j,k \in i} p_{j,k} \bar{P}_j q_{j,k}}{\bar{X}_i} \\ &= \frac{\sum_{j,k \in i} p_{j,k} \bar{X}_{i,j,k}}{\bar{X}_i} \\ &= \sum_{j,k \in i} p_{j,k} \frac{\bar{X}_{i,j,k}}{\bar{X}_i} \\ &= \sum_{j,k \in i} p_{j,k} \omega_{i,j,k}, \end{aligned}$$

where the second line follows from (A.53) and (A.54), the third line follows from (A.57) and (A.2), the fourth line follows from (A.55), and the last line follows from (A.58). We then substitute (A.10) into (A.60) and obtain the following identity:

$$\begin{aligned}
 \text{(A.61)} \quad p_i &= \sum_{j,k \in i} [\mu_j + (\mu_s - \mu_j) + (\mu_{j,s} - \mu_s) + (p_{j,k} - \mu_{j,s})] \omega_{i,j,k} \\
 &= 1 + \sum_{j,k \in i} (\mu_s - \mu_j) \omega_{i,j,k} + \sum_{j,k \in i} (\mu_{j,s} - \mu_s) \omega_{i,j,k} + \sum_{j,k \in i} (p_{j,k} - \mu_{j,s}) \omega_{i,j,k}.
 \end{aligned}$$

We decompose p_i by taking variances across households of both sides of (A.61).

REFERENCES

- AGUIAR, M., AND E. HURST, "Consumption versus Expenditure," *Journal of Political Economy* 113(5) (2005), 919–48.
- and ———, "Life-Cycle Prices and Production," *American Economic Review* 97(5) (2007), 1533–59.
- AGUIRREGABIRIA, V., "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies* 66(2) (1999), 275–308.
- ALBRECHT, J., F. POSTEL-VINAY, AND S. VROMAN, "An Equilibrium Search Model of Synchronized Sales," *International Economic Review* 54(2) (2013), 473–93.
- BAYE, M. R., J. MORGAN, AND P. SCHOLTEN, "Information, Search, and Price Dispersion," in T. Hendershott, ed., *Handbook on Economics and Information Systems*, Vol. 1 (Amsterdam: Elsevier, 2006), 323–77.
- BRODA, C., AND J. PARKER, "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," Mimeo, MIT, 2012.
- BRONNENBERG, B., J. DUBE, M. GENTZKOW, AND J. SHAPIRO, "Do Pharmacists Buy Bayer? Sophisticated Shoppers and the Brand Premium," Mimeo, University of Chicago, 2013.
- BURDETT, K., AND K. L. JUDD, "Equilibrium Price Dispersion," *Econometrica* 51 (1983), 955–69.
- BUTTERS, G. R., "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies* 44(3) (1977), 465–91.
- CHU, C. S., P. LESLIE, AND A. SORENSEN, "Bundle-Size Pricing as an Approximation to Mixed Bundling," *American Economic Review* 101 (2011), 263–303.
- CONLISK, J., E. GERSTNER, AND J. SOBEL, "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics* 99(3) (1984), 489–505.
- EDEEN, B., "Price Dispersion and Demand Uncertainty: Evidence from US Scanner Data," Technical Report, Vanderbilt University Department of Economics, 2013.
- EINAV, L., E. LEIBTAG, AND A. NÉVO, "Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?" *Quantitative Marketing Economics* 8(2) (2010), 207–39.
- GALENIANOS, M., PACULA, AND N. PERSICO, "A Search-Theoretic Model of the Retail Market for Illicit Drugs," *Review of Economic Studies* 79(3) (2012), 1239–69.
- GOLOSOV, M. AND R. E. LUCAS, "Menu Costs and Phillips Curves," *Journal of Political Economy* 115(2) (2007), 171–99.
- HANDBURY, J., "Are Poor Cities Cheap for Everyone? Non-homotheticity and the Cost of Living across US Cities," Mimeo, University of Pennsylvania, 2012.
- HONG, H., AND M. SHUM, "Using Price Distributions to Estimate Search Costs," *RAND Journal of Economics* 37(2) (2006), 257–75.
- KAPLAN, G., AND G. MENZIO, "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," NBER Working Paper, 2014.
- KLENOW, P. J., AND B. A. MALIN, "Microeconomic Evidence on Price-Setting," in B. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, Vol. 3 (Amsterdam: Elsevier, 2010), 231–84.
- McAFEE, R. P., "Multiproduct Equilibrium Price Dispersion," *Journal of Economic Theory* 67(1) (1995), 83–105.
- , J. McMILLAN, AND M. D. WHINSTON, "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quarterly Journal of Economics*, 104 (1989), 371–383.
- MORAGA-GONZÁLEZ, J. L., AND M. R. WILDENBEEST, "Maximum Likelihood Estimation of Search Costs," *European Economic Review* 52(5) (2008), 820–48.
- SOBEL, J., "The Timing of Sales," *Review of Economic Studies* 51(3) (1984), 353–68.
- SORENSEN, A. T., "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy* 108(4) (2000), 833–50.
- VARIAN, H. R., "A Model of Sales," *American Economic Review* 70(4) (1980), 651–9.
- WOLINSKY, A., "True Monopolistic Competition as a Result of Imperfect Information," *Quarterly Journal of Economics* (1986), 493–511.
- WOODWARD, S. E., AND R. E. HALL, "Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage-Market Evidence," *American Economic Review* 102(7) (2012), 3249–76.
- ZHOU, J., "Multiproduct Search and the Joint Search Effect," *American Economic Review*, 104 (2014), 2918–39.

