Store expensiveness and consumer saving: Insights from a new decomposition of price dispersion*

Sofronis Clerides† Pascal Courty‡ Yupei Ma§

January 2022

Abstract We build on recent work analyzing consumers' ability to save by exploiting price dispersion in grocery stores. We show that store expensiveness is not universal but varies across consumers depending on the basket they consume. We incorporate this insight into a price variance decomposition that is an adaptation of Kaplan and Menzio's (2015) approach. We find that the ability to choose the right product at the right store is much less important than Kaplan and Menzio found; rather, the ability to choose the cheapest stores for one's basket is a more important source of variation in consumer savings. Our approach also provides an informal test for competing theories modeling consumers as either shopping for products or shopping for categories, and finds support for both.

Keywords: Price dispersion, grocery shopping, consumer saving, store expensiveness, consumer basket.

JEL Classification: D12, D14.

1 Introduction

There is considerable evidence of substantial dispersion in the prices of grocery store goods. Prices for identical products vary across stores at any given point in time, and across time in any given store. Some stores are cheaper than other stores overall, but not all products are cheaper in those stores. In principle, price-sensitive consumers can exploit this variation in order to purchase their desired basket of goods at a lower overall cost. Does this happen in practice? Can consumers achieve significant savings simply by shopping from cheap stores? Or do they need to engage in time-consuming price comparisons in order to fully exploit the available saving potential?

*We thank Marina Antoniou and Ružica Savčić for research assistance and Christis Tombazos for helpful comments. We gratefully acknowledge funding by the University of Cyprus (internal research grant "Consumer Planning" and Mitacs through its Globalink Research Internship program.

†University of Cyprus and CEPR; s.clerides@ucy.ac.cy (corresponding author).

‡University of Victoria and CEPR; pcourty@uvic.ca.

§University of British Columbia; yupei.ma@alumni.ubc.ca.

The intertemporal dimension of saving has been explored in a literature going back at least twenty years in economics and even further in marketing.[1] The basic idea is that temporary price promotions are an instrument of intertemporal price discrimination between consumers with varying tendencies to optimize over time. The main finding is that consumers who do optimize can achieve substantial savings compared to those who do not. A more recent literature has focused on the multi-product and multi-store nature of grocery shopping. Consumers shop from multiple stores, choosing some products from one store and other products from another. This allows consumers with low search costs to save relative to those who shop from a single store or do not compare prices across stores.

Two papers have used detailed consumer level data to measure the relative importance of the different channels of saving. Griffith, Leibtag, Leicester, and Nevo (2009) explore four ways in which consumers can save: by buying on sale; in bulk; generic; and from low-price outlets. Using data from the UK, they calculate the amount each household saves from each saving channel relative to a benchmark "full" price. They conclude that "the average consumer realizes significant savings from the four dimensions of choice that we study, and that the savings are comparable in magnitude."[2]

Kaplan and Menzio (2015, henceforth KM) focus on heterogeneity in consumer saving. They compute a household price index and decompose it into three components: the *store component*, which captures the overall expensiveness of stores visited by the consumer; the *store-specific good component* (we call it *store-good component* for short), a measure of the expensiveness of products purchased by the consumer relative to the store's overall expensiveness; and the *transaction component*, which represents the expensiveness of products purchased at a store relative to the products' average price in the store. KM's key finding of interest is that the store-good component is more important than the store component in explaining the variance in the prices paid by households; about 50% of the variance is due to the store-good component, 40% to the store component, and only 10% is due to the transaction component.[3] They conclude that there seems 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".[4] KM sketch a model explaining this variation as a result of heterogeneous search behavior: some consumers compare prices across stores, while others

[1] A more detailed discussion of the literature is provided in section 2.

[2] Griffith, Leibtag, Leicester, and Nevo (2009), p. 100.

[3] KM mistakenly reverse these percentages when they discuss these results in their introduction: they attribute 50% to the store component and 40% to the store-good component. This small error does not impact their conclusions.

[4] Kaplan and Menzio (2015), p. 25.

do not.

Our contribution to this literature is based on the simple insight that store expensiveness is not universal, but may differ across consumers depending on the basket they consume. In other words, one consumer's typical basket may be cheaper in store A while another consumer's basket may be cheaper in store B. This is indeed the case in our data: one quarter of consumer baskets cost less in a store that is *more* expensive according to a general price index. Store expensiveness is therefore basket-dependent and consumer-specific. This observation provides the motivation for an alternative decomposition of variance that allows for the possibility that consumers choose the best store for their basket. This amounts to introducing an additional component to the KM decomposition based on our idea of a basket-specific measure of store expensiveness. The resulting extended decomposition (which we humbly call the CCM decomposition) effectively breaks up the KM store-good component into two parts, the *pure store-good component* and the *store-basket component*. The pure store-good component measures the expensiveness of products purchased by the consumer relative to the store's basket-specific expensiveness. We call this the *pure* store-good component because it measures variation in consumer ability to choose the right product in the right store, where the right store is determined on the basis of the cost of the consumer's basket rather than the overall market basket. The store-basket component measures the expensiveness of the consumer's basket at the stores she visits relative to the average expensiveness of those stores. It captures variation in the extent to which the cost of a consumer's basket is representative of the expensiveness of the average store she visits.

We applied both the KM and the CCM decompositions to the IRI Marketing Data Set, which is similar to the Kilts-Nielsen Consumer Panel used by KM. In the main part of our analysis we suppress the transaction component in order to focus on the store and store-good components. With the KM decomposition, the store-good component explains 55% of the variance of the household price index while the store component accounts for 49%. In the CCM decomposition, the pure store-good component explains a modest 24% of the variance, a bit less than half what the KM store-good component explains. Our new store-basket component comes in as a substantial contributor to variance at 46%, while the store component is the largest contributor with 49%.[5] Based on these findings, we conclude that the degree to which households differ in their ability to capture price differences for the goods they purchase at the stores they visit is not as large as KM found. Rather, households differ more in their ability to save by selecting stores that are cheap *for the basket they purchase*.

We obtain further insights by analyzing each product category separately. In the KM model

[5]The variance terms do not add up to 100% because of the covariance terms (see Table 6).

of consumer shopping (developed further in Kaplan, Menzio, Rudanko, and Trachter, 2019, henceforth KMRT), shoppers compare prices of all products across stores. Thomassen, Smith, Seiler, and Schiraldi (2017, henceforth TSSS) use an alternative search protocol where consumers concentrate expenditure for each product category in a single store. The two protocols have different implications for the size of the components at the category level. If all consumers adopt the KM/KMRT protocol, the size of the components at the category level should be similar to the size of the aggregate ones. If consumers follow the TSSS protocol, there should be no pure store-good component, since consumers do not shop around within category.[6]

When we apply the CCM decomposition separately to each of our five product categories, a similar pattern emerges: the pure store-good component is lower than in the aggregate decomposition, while the store-basket component is higher and the store components lower for most categories. The decrease in the role of the store-good component across the board is consistent with the TSSS search protocol, where consumers shop for categories rather than individual goods. But the fact that the pure store-good component is not zero suggests that some consumers do compare prices of the same good across stores, as in the KM/KMRT model.

An additional contribution of our work is that it elucidates the inner workings of KM's decomposition methodology. Our exposition of the methodology uses an alternative formulation based on hypothetical price indexes that correspond to different shopping protocols. For example, we constructed the store-good hypothetical price index, which is the cost of the consumer's basket had she bought each item at the average price of the store she purchased it from. The decomposition is then defined as the sum of differences between pairs of price indexes, which represent differences in the cost of the consumer's basket under different protocols. The formulation in terms of price indexes allows for intuitive interpretations and provides useful economic insights. It is also flexible and general, as it is easy to define new price indexes in order to analyze different dimensions of heterogeneity in consumer saving. We believe that our exposition helps make the methodology more transparent and accessible. Our analysis also provides support for the usefulness of the methodology. We find that most of the KM findings carry through to a different data set and are robust to a variety of different assumptions. An important exception is the transaction component, which is significantly higher in our data than in the KM data. This finding survived a barrage of robustness tests and remains a puzzle for further exploration.

The rest of the paper is organized as follows. Section 2 briefly summarizes the literature on

[6] These implications derive formally from Proposition 1 in section 4.

price dispersion and consumer saving. Section 3 explains the data and provides descriptive statistics. Section 4 presents the variance decomposition methodology and uses a simple example to illustrate our notion of the store-basket price index and to highlight the differences between the KM and CCM approaches. The main results from applying both methodologies are presented in Section 5, along with several robustness tests. Section 6 concludes.

2 Literature

In an early contribution, Pratt, Wise, and Zeckhauser (1979) noted the existence of different prices in markets they describe as 'almost competitive'. Systematic evidence of price dispersion began to accumulate in the 2000s with studies of small numbers of products (Sorensen, 2000; Lach, 2002). The increased availability of large and detailed datasets in the last few years has made it possible to study price dispersion using thousands or even millions of products. The grocery sector has been the subject of many of these studies, such as Hosken and Reiffen (2004), Kaplan and Menzio (2015), Dubois and Perrone (2019), Moen, Wulfsberg, and Aas (2020), and Hitsch, Hortaçsu, and Lin (2021). They all document large and persistent price dispersion for narrowly defined products sold in grocery stores.

A different strand of the literature focused on intertemporal price variation in the form of sales promotions of specific products. The review article by Neslin (2002) is a good source for the large marketing literature on this topic. Pesendorfer (2002) was an early contribution to the economics literature. Dynamic inventory models for the problem of intertemporal optimization of storable good purchases were later developed by Erdem, Imai, and Keane (2003), Hendel and Nevo (2006a) and Hendel and Nevo (2013). Seiler (2013) and Pires (2016) developed inventory models that incorporated the decision to engage in costly search. Clerides and Courty (2017) showed that consumers often miss opportunities to buy cheap – even in cases when the search cost appears minuscule – and attribute this behavior to inattention. The emphasis in this literature is on price comparisons over time and/or across brands, but not across stores within the same time period.

DellaVigna and Gentzkow (2019) have shown that U.S. chains in a broad range of retail sectors charge nearly uniform prices across their stores. In the grocery sector, Hitsch, Hortaçsu, and Lin (2021) have shown that prices vary across stores within the same market but less so across stores within the same retail chain. In other words, prices are set at the chain level and do not adjust to local conditions. KMRT document that a significant source of price dispersion across stores is due to persistent differences in the price that different retailers set for a good relative to the price they set for other goods; they call this type of price variation *relative price dispersion*. KMRT develop a model that delivers relative price dispersion as an equilibrium outcome. Sellers in the model are homogeneous while buyers are heterogeneous. One type of buyer (the 'busy' type) has a high valuation for the goods and needs to purchase all the goods at the same location. The other type of buyer (the 'cool' type) has a low valuation for the goods and is able to purchase different goods at different locations. The cool consumers have low valuations for the goods and therefore have an incentive to search across stores for low prices.

KMRT build on a theoretical literature of price search dating back at least to Varian's (1980) classic model of sales. A notable example of the more recent empirical literature is TSSS, who develop a multi-category, multi-seller demand model and estimate it using grocery store data from the UK. Stores in their model sell different categories of products, such as household goods, drinks, fruits and vegetables, meat, etc. Consumers select one store for each category; that is, they shop around for categories rather than for individual goods. Some consumers tend to shop in a single store; the existence of these consumers is important because they generate relatively large cross-category effects and therefore have a greater pro-competitive impact.

3 Descriptives

3.1 Data

We use the well-known IRI Marketing Data Set.[7] The dataset provides store level sales and price information for 30 product categories in 47 U.S. markets over the 12-year period 2001-2012. For two of those markets (Eau Claire, Wisconsin and Pittsfield, Massachusetts) additional data on consumer purchases are available through the Behavior Scan panel. A total of about ten thousand distinct households are represented on the panel, with an average of roughly five thousand households every year. Behavior Scan includes information on every shopping trip made by each participating panelist during the sample period.[8] For each trip, it records the number of units purchased of each good (defined as a unique UPC) and the unit price. The raw dataset also records the number of stores visited by each panelist in each quarter.[9]

[7]See Bronnenberg, Kruger, and Mela (2009). The dataset has been widely used in this literature, including recently by Pires (2016) and Ching and Osborne (2020). It is similar in structure and content as the *Nielsen* dataset used by KM, though it is not as extensive.

[8]We use the terms panelist, consumer and household interchangeably.

[9]Appendix C replicates the KM decompositions using the IRI sample. Most of KM's findings carry through with one exception (see Section 6).

We work with the top five product categories in terms of total purchase count: carbonated soft drinks, milk, salty snacks, yogurt, and cold cereal. The sixth and seventh categories (soup and frozen dinners) could not be used because they had missing product characteristic values that prevented us from accurately sorting UPCs into unique products. The five categories selected cover 55% of all purchases in the dataset; adding a few more product categories would only marginally increase this figure. The online data appendix explains how we merged UPCs into unique products and how we removed products, stores and panelists with few purchases.[10]

Table 1 provides some summary information about the panelists, products and purchases in our final sample. Panelists stay on average about six years in Behavior Scan. Each quarter, they visit on average 2.2 stores, buy eighteen distinct products from four of the five categories, and make close to thirty purchases total. The summary statistics are broadly similar in Eau Claire and Pittsfield. The most notable difference is in the number of stores visited: Pittsfield residents visit 2.41 different stores per quarter, versus 2.00 for Eau Claire residents.

Table 1: Information about the final IRI sample

	Eau Claire	Pittsfield
Observation count		
Quarters	48	48
Goods	3,812	3,836
Purchases	3,862,540	3,977,461
Panelists	5,609	5,144
Stores	6	7
Averages across panelists		
# quarters panelists remain in the dataset	23.31	24.98
# distinct goods bought per quarter	17.82	18.49
# categories purchased per quarter	3.94	3.96
# stores visited per quarter	2.00	2.41
# purchases per quarter	26.94	29.04

3.2 Store visits, product availability, and price comparisons

Consumers' ability to save from cross-store comparison shopping depends on the number of stores they visit and on the availability of products across the different stores. We say that a product is available in a given store and quarter if the store records a positive quantity for

[10]Note to referees: the online appendix has been submitted together with the paper.

that product-quarter (see online appendix). This section establishes some important stylized facts about store visits and product availability.

1. *Although store availability varies greatly across products, a significant share of products is available in all stores.*

   For each market-quarter pair, we counted the number of stores each product was sold in. For reference, Table 1 reports that there are 6 stores in Eau Claire and 7 in Pittsfield. The median product was available in 5 stores in both towns; 27.4% of all products are available in all stores of a market. Only 9.3% of products are available in a single store, and this is similar in the two markets. This figure drops to 2.66% when we compute availability at the transaction instead of product level.[11] Overall, there is substantial variation in product availability across markets (lower in Pittsfield than in Eau Claire), product categories (lower for milk and yogurt, higher for carbonated soft drinks) and product popularity (higher for products with larger market shares).

2. *Consumers visit few stores and do most of their spending in their top two stores.*

   We counted the number of stores visited for each consumer-quarter pair. We find that 27.9% of consumers visit a single store in a given quarter and 84.1% visit at most 3 stores. We computed for every consumer-quarter the fraction of the consumer's expenditure spent in each of the top two stores she does most of her shopping at. On average, consumers do 77.3% of their spending in the single store they frequent most, and 94.1% in two stores.[12]

3. *Most purchased products are available in most visited stores.*

   For each product purchased, we computed the fraction of stores in which the product was available among those stores visited by the consumer in the same quarter. The average of this fraction over all purchases is 90.9%; a full 80% of transactions are available in all stores visited. The figures suggest that, for the large majority of instances, consumers had an opportunity to purchase the same product in another store they visited in the same quarter. There is little variation across consumers with respect to this finding. The vast majority of consumers (95%) can find the majority of the products they purchase (74.8%) in all stores they visit.[13]

[11] The difference with the product level computation is that at the transaction level a product that is purchased many times will be counted every time, as opposed to just once per market-quarter.

[12] TSSS report very similar figures for their UK data: 71% and 94%.

[13] Another way to measure the extend to which store unavailability prevents price comparison, it to use the notion of pairwise price comparison. A pairwise price comparison is possible for a purchased product and a store visited different from the one where the product was purchased, if the product is available in that

4. *The main reason some consumers cannot compare prices across stores is that they visit a single store.*

   For 32.7% of transactions, the consumer cannot make a price comparison. In most cases (73.0%), this occurred because the consumer visited a single store. Among the consumers who visit multiples stores, price comparisons are possible for 88.5% of transactions. One price comparison is possible in 42.9% of transactions, two in about 27.4% and more than two in about 18.2%.

We conclude that product availability in not a major impediment to price comparisons: consumers can find the majority of the products they purchase in most of the stores they visit. The reason why price comparison is not possible for about a third of the transactions is that almost one third of consumers visit a single store.

3.3 Household price indexes

A key part of our analysis is the construction of household price indexes (HPIs). The *actual* HPI is defined as in Aguiar and Hurst (2007) and Kaplan and Menzio (2015) as the ratio of the actual cost of a consumer's basket to the cost of the same basket at the average market price of each product. In addition, we define *hypothetical* HPIs that give us the cost of the consumer's basket under different shopping scenarios. The calculations for the indexes are fairly complex and explaining them in full detail would require some cumbersome notation. To simplify things as much as possible, in the exposition below we show how to construct the HPIs using consumer purchases in a single market and a single period (set to a quarter). We use the term HPI, omitting the qualifier actual or hypothetical, when this is obvious from the context.

With these simplifications, an observation is indexed by $i = 1 \ldots I$ for panelist, $j = 1 \ldots J$ for good, and $s = 1 \ldots S$ for store. Two variables contain all relevant information: $q_{i,j,s}$ is the number of units of good j purchased by individual i at store s; and $P_{i,j,s}$ is the average unit price paid. With temporal price variation, the average unit price may vary across households, and this could be due to (i) chance (some households happen to purchase when the price is low, others when the price is high) or (ii) household heterogeneity (bargain hunters versus loyal consumers in price discrimination literature).

Table 2 explains how the household price indexes are computed. In the first step (section A

other store. The ratio of all possible pairwise price comparisons, to the maximum number of possible pairwise price comparisons, were purchased products available in all stores visited, is 80.2%. This demonstrates that product availability does not prevent consumers from comparing prices.

Table 2: Normalized prices, expenditure shares, and household price indexes

A. Weighted average normalized prices		
Normalized price	KM1	$\mu_{i,j,s} = \frac{P_{i,j,s}}{P_j}$
Average market price	KM2	$P_j = \sum_{i,s} P_{i,j,s} \frac{q_{i,j,s}}{\sum_{i,s} q_{i,j,s}}$
Weighted average normalized prices:		
Market	KM3	$\mu_j = \sum_{i,s} \mu_{i,j,s} \frac{q_{i,j,s}}{\sum_{i,s} q_{i,j,s}} = 1$
Good j in store s	KM4	$\mu_{j,s} = \sum_{i} \mu_{i,j,s} \frac{q_{i,j,s}}{\sum_{i} q_{i,j,s}}$
Store s	KM5	$\mu_s = \sum_{j} \mu_{j,s} \frac{\sum_i P_{i,j,s} q_{i,j,s}}{\sum_{i,j} P_{i,j,s} q_{i,j,s}}$
B. Household expenditure shares		
Share of expenditure on j in s	KM13	$\omega_{i,j,s} = \frac{P_j q_{i,j,s}}{\sum_{j,s} P_j q_{i,j,s}}$
C. Household price indexes		
Actual	KM12	$p_i = \sum_{j,s} \mu_{i,j,s} \omega_{i,j,s} = \frac{\sum_{j,s} P_{i,j,s} q_{i,j,s}}{\sum_{j,s} P_j q_{i,j,s}}$
Market		$p_i^m = \sum_{j,s} \mu_j \omega_{i,j,s} = 1$
Store-good		$p_i^{sg} = \sum_{j,s} \mu_{j,s} \omega_{i,j,s}$
Store		$p_i^s = \sum_{s} \mu_s \omega_{i,\cdot,s}$
Store-specific basket*		$p_{i,s} = \sum_{j} \mu_{j,s} \omega_{i,j,\cdot}$
Store-basket		$p_i^{sb} = \sum_{s} p_{i,s} \omega_{i,\cdot,s} = \sum_{j,s} \mu_{j,s} \omega_{i,\cdot,s} \omega_{i,j,\cdot}$

Notes: (1) The second column provides the reference to the equation number in Kaplan and Menzio (2015), although we have adapted some notations for the sake of clarity. (2) A dot '.' in a variable's subindex means that the variable is summed over that subindex, i.e. $\omega_{i,\cdot,s} = \sum_j \omega_{i,j,s}$. (∗) Appendix A explains how store-specific baskets are reweighed to account for partial product availability.

in the table) we divide each price $P_{i,j,s}$ by the average market price P_j to obtain $\mu_{i,j,s}$, the normalized average price paid by panelist i for good j in store s. In the second step (also in section A) we use the $\mu_{i,j,s}$'s to compute three weighted average normalized prices: μ_j, the average market price that is equal to unity by definition; $\mu_{j,s}$, the average price of good j in store s; and μ_s, which is the price level of store s relative to the overall price level in the market.

The third step computes the desired price indexes by taking weighted averages of the μ's, with the weights $\omega_{i,j,s}$ being the expenditure shares of each household on product j in store s (as a fraction of total household expenditure – section B of Table 2). The first index defined in section C of the table is p_i, the actual HPI used by KM, and it is computed using the normalized price paid by the consumer for her basket. Equivalently, p_i can be computed (as in KM) as the ratio of the actual cost of household i's shopping basket, $\sum_{j,s} P_{i,j,s}q_{i,j,s}$, and the cost of the same basket had the household paid the average market price for each item, $\sum_{j,s} P_j q_{i,j,s}$. As such, this index is a measure of the household's ability to save: a household saves when $p_i < 1$, meaning that it disproportionately purchases products with lower relative prices. The household dissaves when $p_i > 1$. There is no saving or dissaving on average across all households: the average HPI, using household expenditure as weights, is equal to one.[14] For the sake of brevity, we will be talking about household saving, keeping in mind that all statements equally apply to dissaving.

The other price indexes in section C of Table 2 are alternative ways of calculating the cost of the consumer's basket using hypothetical prices rather than actual ones:

p_i^m is the cost of the consumer's basket had she bought each item at the average market price, and is equal to unity by definition;

p_i^{sg} is the cost of the consumer's basket had she bought each item at the average price of the store she purchased it from;

p_i^s is the cost of the consumer's basket on the basis of the average expensiveness of the stores she purchases each item from; put differently, it is the average expensiveness of the stores the consumer visits, evaluated on the basis of her basket;

p_i^{sb} is the cost of the consumer's basket had she purchased all items in her basket in each of the stores she visits in proportion to her overall spending in each store; it measures the expensiveness of the panelist's basket at the stores visited.

These indexes are not explicitly defined in KM. The first three appear implicitly in KM

[14] $\sum_i \alpha_i p_i = 1$ for $\alpha_i = \frac{\sum_{j,s} P_j q_{i,j,s}}{\sum_{i,j,s} P_j q_{i,j,s}}$.

Table 3: Average statistics of price and HPIs

	Nielsen (KM)		IRI				
	Price	p_i	Price	p_i	p^{sg}	p^s	p^{sb}
Std dev.	0.21	0.09	0.16	0.07	0.04	0.03	0.04
90-10 ratio	1.79	1.22	1.54	1.18	1.10	1.07	1.09
90-50 ratio	1.29	1.09	1.20	1.08	1.05	1.04	1.04
50-10 ratio	1.38	1.12	1.29	1.09	1.05	1.04	1.04

The two columns under the heading *Nielsen (KM)* are copied from KM. The next two columns provide the same statistics with our IRI data, and the last three give statistics on the hypothetical price indexes.

equation (14), while the fourth one is the new index introduced in this paper in order to explore the idea of consumer-specific store expensiveness.

Table 3 reports distributional statistics on the HPIs. The two columns under the heading *Nielsen (KM)* are copied from KM (Table 2, column 2 on page 9 and Table 6, column 1 on page 22 respectively). They report expenditure-weighted averages across markets and quarters of measures of dispersion for the price $\mu_{i,j,s}$ and the actual HPI p_i. The 90-10 ratio is the ratio of the price at the 90th percentile to the price at the 10th percentile; the other ratios are defined in a similar way. The next two columns report the statistics for the same variables in our IRI dataset.

The variation in our data is somewhat smaller – the average standard deviation is 21% with the Nielsen data and 15% with the IRI data – but the overall patterns are similar. Some panelists spend a significantly larger amount on grocery relative to others: in our IRI data, the panelist at the 90th percentile spends 18% more on groceries than the panelist at the 10th percentile. The equivalent figure for the Nielsen data is 22%. The last three columns report the same statistics for the store-good, store, and store-basket indexes respectively. The standard deviations of the store-good and store-basket price indexes are greater than the standard deviation of the store price index.

4 Decomposition

4.1 Decompositions in terms of price indexes

The KM decomposition of the HPI (equation 14) can be written in terms of the price indexes as follows:

$$p_i = p_i^m + \underbrace{(p_i - p_i^{sg})}_{\text{transaction}} + \underbrace{(p_i^{sg} - p_i^s)}_{\text{KM store-good}} + \underbrace{(p_i^s - p_i^m)}_{\text{store}}. \quad (1)$$

Casting the KM decomposition in terms of price indices allows for an intuitive interpretation of the components. Each price index gives the cost of the consumer's basket under a specific shopping plan. By comparing price indexes, we can calculate the savings the consumer can make by adopting one shopping plan over another. For example, the difference $p_i - p_i^{sg}$ (KM's transaction component) tells us the cost of the consumer's basket relative to its cost had she paid the average store price for each item. Therefore the difference tells us how much she (dis)saved by timing her purchases.

Since our emphasis is on the store-good component $p_i^{sg} - p_i^s$, we will suppress the transaction component for the main part of the analysis.[15] We rewrite equation (1) as

$$p_i^{sg} - p_i^m = \underbrace{(p_i^{sg} - p_i^s)}_{\text{KM store-good}} + \underbrace{(p_i^s - p_i^m)}_{\text{store}}, \quad (2)$$

The left-hand side is the difference between the cost of the consumer's basket at the average price of the store she purchased it from and the cost of the basket at the average market price. This difference is attributed to the store and KM store-good components. The store component measures the expensiveness of the stores visited by the consumer relative to the market. It therefore tells us how much the consumer (dis)saves by shopping in the chosen stores. The KM store-good component is the expensiveness of the household's basket in those stores relative to overall store expensiveness. It measures how much the consumer (dis)saves by selecting the right product (in terms of price) from the right store among the stores visited.

We argue that, if our objective is to assess the household's ability to choose cheap products, then p_i^s is not the best benchmark to compare p_i^{sg} to. The reason is that p_i^s is based on a measure of store expensiveness, μ_s, that is calculated on the basis of *all* products. It is more appropriate to compare p_i^{sg} to our proposed new index p_i^{sb}, which is calculated using a

[15] We bring back the transaction component when we discuss robustness in section 5.4.

household-specific measure of store expensiveness, $p_{i,s}$. Thus, our innovation is to introduce p_i^{sb} and use it to define a finer decomposition of $p_i^{sg} - p_i^m$:

$$p_i^{sg} - p_i^m = \underbrace{(p_i^{sg} - p_i^{sb})}_{\text{pure store-good}} + \underbrace{(p_i^{sb} - p_i^{s})}_{\text{store-basket}} + \underbrace{(p_i^{s} - p_i^{m})}_{\text{store}}, \tag{3}$$

Essentially, we have broken down the KM store-good component into two, the *pure store-good component* and the *store-basket component*. The former is our measure of the household's ability to choose cheap products from the stores it visits. The latter measures the extent to which the household purchases a basket that is representative of the expensiveness of the stores it visits.

The store-good component is zero in three benchmark cases worth discussing. Proposition 1 formerly describes the three cases (see Appendix B for the proof).

Proposition 1. *A consumer has zero store-good savings, $p_i^{sb} = p_i^{sg}$, when: (a) she visits a single store; (b) store-good prices $\mu_{j,s}$ do not vary across stores visited; or (c) she purchases the same share of each good in all stores visited ($\frac{\omega_{i,j,s}}{\sum_j \omega_{i,j,s}}$ constant across s).*

Although the proposition only states sufficient conditions, it highlights benchmark cases discussed in the literature. Statement (a) says that consumers who visit a single store (about one third of consumers in our sample – see section 3.2) cannot save by comparing prices across stores. This is important because the KM decomposition incorrectly attributes the savings of these consumers to the store-good component (see our example in section 4.2). An illustration of part (b) is a consumer who is loyal to a retail chain. This is relevant because the literature has shown that uniform pricing tends to hold within a chain, but not for stores that belong to different chains (DellaVigna and Gentzkow, 2019; Moen, Wulfsberg, and Aas, 2020; Hitsch, Hortaçsu, and Lin, 2021). Again, such a consumer will have a zero pure store-good component. Statement (c) is an illustration of the 'naive' or 'busy' consumers in the price discrimination literature (Lal and Matutes, 1989; Kaplan, Menzio, Rudanko, and Trachter, 2019). These consumers have high search costs and do not take the time to compare relative prices at the stores they visit; they visit multiple stores but the baskets they purchase from each store are composed of the same goods purchased in the same proportions. To summarize, a high store-good component requires a consumer to visit multiple stores, prices to vary across stores visited, and the consumer to systematically collect and compare the different prices.

The conditions stated in Proposition 1 are influenced by both consumer and store behavior, in the sense that the attribution of consumer savings to the store-basket or store-good component depends on the number of stores visited and on store pricing and product assortment policies. To illustrate, consider a simplified market where: (a) products are sold at normalized prices that vary across products and stores and can take only one of two values, $\mu_{j,s} = c$ for cheap or $\mu_{j,s} = e$ for expensive, and (b) all stores sell the same expenditure share of expensive products.[16] This latter assumption implies $\mu_s = 1$ for all stores and there is no store component, $p^s = 1$. With this as background, we now consider a consumer who buys n cheap products from n different stores. We have $p^{sg} = c$ and the consumer savings are $p^s - p^{sg} = 1 - c$. In one scenario, each product is cheap in only one of the visited stores. If the consumer spends the same amount on each product, we obtain that her store-specific basket (see Table 2 Panel B) is composed of $\frac{1}{n}$ and $\frac{n-1}{n}$ of cheap and expensive products respectively in any of the stores she visits, and $p^{sb} = \frac{1}{n}c + \frac{n-1}{n}e$. The store good component, $p^{sb} - p^{sg} = \frac{n-1}{n}(e - c)$, increases as the consumer visits more stores. This is as it should be, since buying cheap products requires more cross-store shopping as the number of stores visited increases. In an alternative scenario, where cheap products are cheap in all stores visited, consumer savings are explained by the store-basket component alone, since there is no pure store-good component, $p^{sg} - p^{sb} = 0$. This demonstrates that the attribution of consumer savings to the store-basket and store-good components depends both on consumer behavior (store visited and purchase choices) and store pricing policies (whether store prices are correlated across stores).

The next subsection applies both decompositions in a simple setting that allows us to highlight the differences between them.

4.2 An illustrative example

We have constructed a stylized consumer shopping model that helps clarify the main concepts and highlights the difference between the KM and CCM decompositions. The point of the example is to narrow the analysis down to the KM store-good and CCM pure store-good components: by design, there is no transaction component (the model is static, so there is no intertemporal price variation) and no store component (because stores are perfectly symmetric).

Consider a market with two stores selling the same two products, say bread and milk. One store specializes in bread (call it a bakery) and the other specializes in milk (a dairy). Both stores offer lower prices on their specialty products. Letting $P_{i,j}$ denote the price of good $i \in \{b, m\}$ in store $j \in \{B, D\}$, we have $P_{b,B} < P_{b,D}$ and $P_{m,D} < P_{m,B}$. Consumers differ in

[16] See Appendix B for an example of store pricing policies that generate these normalized prices.

two dimensions: the composition of their basket and their shopping behavior. One fifth of consumers are shoppers and the other four fifths are loyals. Loyals have a basket containing half a unit each of bread and milk. They buy from a single store that could be chosen, say, on the basis of location, and are evenly split across the two stores. Shoppers buy each item in their basket at the lowest available price. They come in three types of equal size: bread shoppers purchase one unit of bread from the bakery, milk shoppers one unit of milk from the dairy, and all-shoppers purchase a half-unit of milk from the dairy and a half-unit of bread from the bakery. This is similar to the Lal and Matutes (1989) setup, where some consumers purchase their entire basket from a single store while others shop around. Table 4 describes the consumer types and their purchases $q_{i,j,s}$.

Table 4: Consumer types and their choices

	(Stores)		Bakery		Dairy	
	(Products)		Bread	Milk	Bread	Milk
	Consumer type	*Frac.*	*Quantity purchased* q_{ijs}			
1	Bread shopper	1/15	1	0	0	0
2	Milk shopper	1/15	0	0	0	1
3	All-shopper	1/15	1/2	0	0	1/2
4	Dairy loyal	2/5	0	0	1/2	1/2
5	Bakery loyal	2/5	1/2	1/2	0	0
	Quantity purchased		3/10	1/5	1/5	3/10

In order to calculate the HPIs we need prices. Suppose $P_{b,B} = P_{m,D} = 1.0$ and $P_{b,D} = P_{m,B} = 1.1$. Table 5 reports the HPIs and the components from the two decompositions with these prices. Since the store component $p^s - p^m$ is zero by construction, the KM decomposition attributes price dispersion for all consumers entirely to the KM store-good component (equation (2)). This is appropriate for the all-shoppers (type 3) because these consumers save by purchasing the goods in their basket from the stores where these goods are cheap. But the nonzero KM store-good component for consumer types 1, 2, 4, and 5 is problematic. For example, consumers 1 and 2 have a negative KM store-good component because the cost of their single-item basket is lower than the overall expensiveness of the store they visit. Yet there is no basis on which to conclude that these consumers buy the right product from the right store, which is the KM interpretation of a nonzero store-good component, as they only purchase a single item from a single store. In contrast, the CCM pure store-good component is zero for all four consumer types (each type satisfies one of the conditions stated in Proposition 1). Consumers 1 and 2 have a negative store-basket

Table 5: Consumer price indexes and decompositions

	Price indexes					Components			
Cons. type	p	p^m	p^s	p^{sg}	p^{sb}	store KM/CCM $p^s - p^m$	store-good KM $p^{sg} - p^s$	pure store-good CCM $p^{sg} - p^{sb}$	store-basket CCM $p^{sb} - p^s$
1	0.96	1.00	1.00	0.96	0.96	0.00	-0.04	0	-0.04
2	0.96	1.00	1.00	0.96	0.96	0.00	-0.04	0	-0.04
3	0.96	1.00	1.00	0.96	1.01	0.00	-0.04	-0.05	0.01
4	1.01	1.00	1.00	1.01	1.01	0.00	0.01	0	0.01
5	1.01	1.00	1.00	1.01	1.01	0.00	0.01	0	0.01
Market level decomposition						0%	100%	39%	72%

Note: The CCM decomposition adds up to 100% once the covariance term, $cov(p^{sb} - p^m, p^{sg} - p^{sb}) = -5.5\%$, is included.

component because they choose the right store for their basket. Consumers 4 and 5, on the other hand, have a small positive store-basket component because their basket contains a larger fraction of the expensive good (in value terms) than the market basket.

The subtle difference between the store-good and store-basket components leads to a striking difference between the two decompositions, as reported on the last line of Table 5. The KM decomposition attributes 100 percent of the dispersion to the KM store-good component while the CCM decomposition attributes only 39 percent to the pure store-good component, with 72 percent being attributed to the store-basket component. The conclusion from the CCM decomposition is that differences in prices consumers pay is primarily due to variation in consumers' ability to select stores on the basis of the expensiveness of their basket in these stores, and less so to the ability of selecting the cheapest products across stores.

4.3 Store expensiveness is basket-specific

In the above example, the store-basket HPIs of consumers 1 and 2 are low relative to their store HPIs ($p^{sb} = .96 < 1 = p^s$). These consumers' baskets are cheap at the stores they shop at, and this is what explains the high store-basket component. To further motivate the CCM decomposition, we present direct evidence that store expensiveness is basket-specific in our sample of households.

For each panelist-quarter observation, we select the top two stores (s_1, s_2) in terms of overall expenditure, and rank them in two ways: according to their basket-specific price level $p_{i,s}$

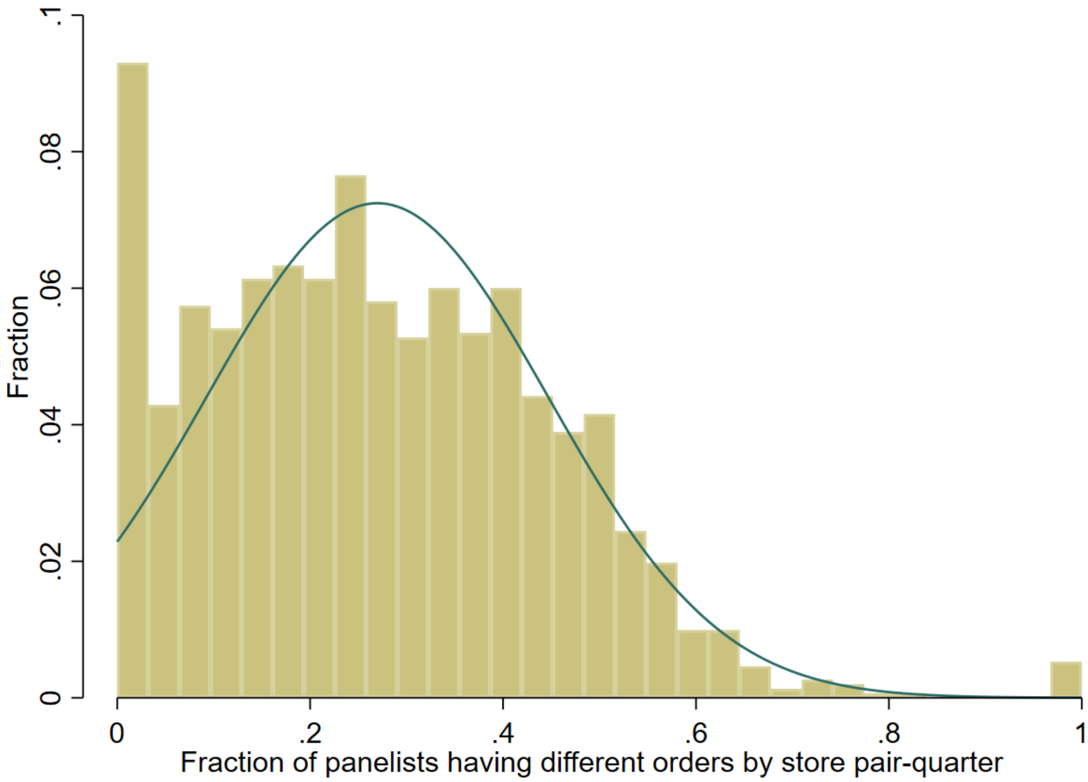

Figure 1: Distribution of disagreement (fraction of panelists with $Sign(p_{i,s_1} - p_{i,s_2}) \neq Sign(\mu_{s_1} - \mu_{s_2})$) by store pair-quarter

and according to their overall price level μ_s. For each store-pair quarter triplet, we compute a measure of disagreement over store ranking defined as the fraction panelists who have different orderings with the two store price indexes, $Sign(p_{i,s_1} - p_{i,s_2}) \neq Sign(\mu_{s_1} - \mu_{s_2})$. Figure 1 plots the distribution of the disagreement measure.[17]

If store expensiveness was not basket-specific, we would expect a large spike as zero. Instead, we find that in only 5% of store-pair quarter observations do households agree (they all rank the two stores the same for both measures).[18] The distribution of disagreement over store ranking has a wide support. On average, store preference is basket-specific for 26 percent of panelists. This evidence supports using a basket-specific measure of store expensiveness as done in the CCM decomposition.

[17]There are 1517 store-pair-quarter observations: both stores in the pair are one of the top two stores by expenditure for at least one panelist in that quarter (the upper bound is 36 store-pairs times 48 quarters = 1728). After filtering out store pair-quarters with fewer than 50 panelists, we end up with 954 observations.

[18]The spike at zero on Figure 1 says that a bit more than 9% of store-pair quarters have 3.2% (bin size of .032) or fewer panelists having different rankings.

5 Results

5.1 Decomposition results

Table 6 shows the results of applying the CCM and KM decompositions to our data. The CCM decomposition is presented in column 2 on the left panel. The KM decomposition is presented in the second to last line of the right panel, where the components are obtained by summing the appropriate terms (marked by "X") in each column. We follow KM in computing statistics for each of the 96 market-quarters (excluding all good-market-quarters with fewer than 25 transactions) and then aggregating them by taking expenditure-weighted averages across markets and quarters.

Table 6: The KM and CCM decompositions

CCM decomposition		KM decomposition		
		store	store-good	cov
var(pure store-good)	24%		X	
var(store-basket)	46%		X	
var(store)	49%	X		
2*covar(store-basket, pure store-good)	-15%		X	
2*covar(store, pure store-good)	4%			X
2*covar(store, store-basket)	-7%			X
sum = var(overall)	100%	49%	55%	-3%
			sum = 100%	

Note: All variances and normalized by the total variance, $var(p^{sg}_{i,t} - p^{m}_{i,t})$. The full CCM decomposition is reported on the left. The KM decomposition is reported on the second to last line and computed as vertical sums of the relevant terms from the CCM decomposition.

Looking first at the CCM decomposition, note that the store-basket component accounts for a bit less than half of the overall variance (46%). This confirms the result from the previous section that store expensiveness and store-basket expensiveness are not the same thing. The store component accounts for about half (49%) of the overall variance. The pure store-good component is the smallest of the three, contributing 24% to the overall variance.

The KM store-good component is more than double the size of the CCM pure store-good component (55% vs 24%). This difference in the estimated store-good components can be seen clearly in Figure 2, which plots their full distributions for an indicative market-quarter.[19] Using the CCM method, there is a large spike at zero: 54% of consumers have a zero pure

[19] Since the variance is computed by market-quarter, it is not possible to show the distributions for all market-quarters in a single graph. The distributions displayed are for Eau Claire 2001:1.

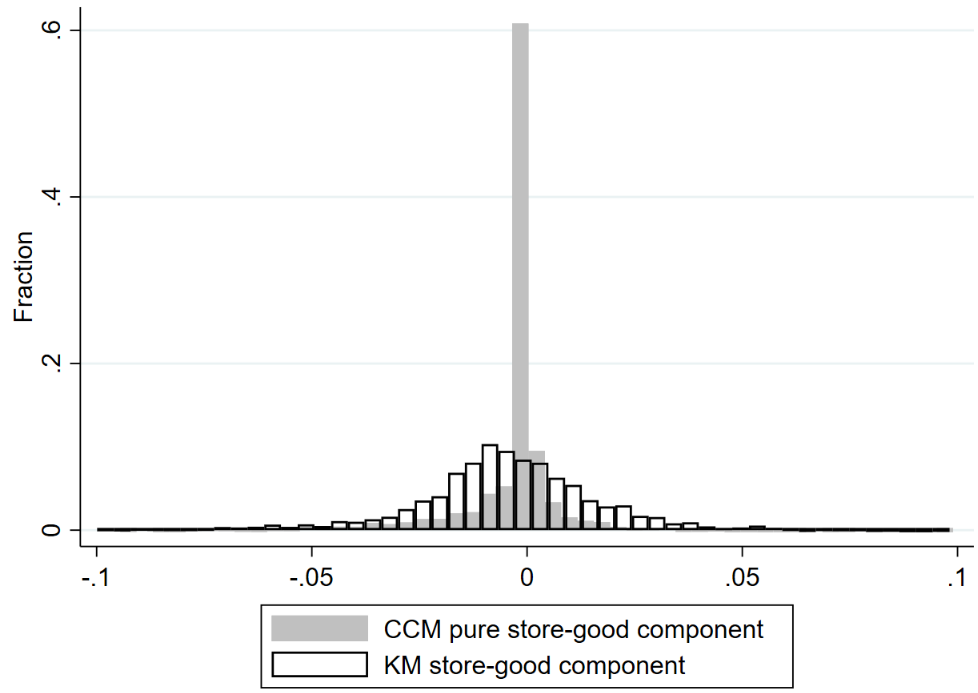

Figure 2: The distribution of the KM store-good component and the CCM pure store-good component in an indicative market-quarter

store-good component.[20] The distribution using the KM method has no such spike and is much more spread out. This is reminiscent of the example reported in Table 5, which showed that the KM store-good component is much larger than the CCM pure store-good component.

5.2 Single-store shoppers

Table 6 includes consumers who visit just a single store (see discussion in section 3.2). These consumers, who make up 27.9% of the total and account for 23.5% of purchases, have no store-good component (Proposition 1). Can the small pure store-good component in the CCM decomposition be explained by the large fraction of households visiting a single store? We computed the CCM decomposition separately for consumers visiting a single store and those visiting multiple stores and report the results in Table 7. The store-good component is higher for multi-store shoppers than for the entire sample (32% versus 24%). Still it is smaller than the store-good component from the KM decomposition (32% versus 55%) and the store-basket component (32% versus 43%). Thus, the conclusion that the store-good

[20]The spike on the figure is slighly higher because the bin size is .0038. Across all markets quarter, 29% of consumers have a zero pure store-good component.

component is small relative to the store-basket component in CCM is not driven by the existence of consumers visiting a single store.

Table 7: The CCM decompositions for households visiting a single and multiple stores

	Entire sample	Multiple stores	Single store
var(pure store-good)	24%	32%	0%
var(store-basket)	46%	43%	52%
var(store)	49%	45%	58%
2*covar(store-basket, pure store-good)	-15%	-19%	0%
2*covar(store, pure store-good)	4%	5%	0%
2*covar(store, store-basket)	-7%	-6%	-10%
sum = var(overall)	100%	100%	100%

Note: All variances and normalized by the total variance, $var(p_{i,t}^{sg} - p_{i,t}^{m})$.

Households visiting a single store have about the same store and store-basket components (58% and 52% respectively). For these households, the store-basket component is attributed to the store-good component under the KM decomposition. The misallocation of the 52% store-basket component to the store-good component for 27.9% of households explains roughly half of the 31% difference in store-good components between the KM and CCM decomposition.

5.3 Decomposition by product category

We can obtain additional insights into consumer behavior by examining each product category separately. To see how, we review two benchmark models of relative price comparison that rely on different search protocols.

Under one view, consumers compare prices of all products across all visited stores, independently of product category. There is a fixed cost of comparing prices for a given product that may differ across consumers but does not depend on the category the product belongs to. We call this the KMRT view because category does not play a role in their theory of relative price dispersion. Under this view, the decomposition should produce the same results when categories are examined together or separately. Under the second view, consumers source all products of a given category from a single store but may source different categories from different stores. This is the search protocol adopted in TSSS. Under this view, the decomposition by category should have a zero pure store-good component (see Proposition 1). In such a scenario, a non-zero aggregate pure store-good component could arise because of variation in ability to correctly choose a store for each category.

The two views outlined above have different implications for how household expenditure shares, and saving decompositions, should change when disaggregating purchases by categories. We now confront these implications to the data, starting with the evidence on store expenditure shares. Recall from Section 3.2 that (across all categories) consumers spend 77.3% of their expenditure in their top store and 16.8% on their second store. Looking at expenditure category by category and taking average across categories, these figures are 86.7% and 12.5% respectively. Although consumers are more likely to concentrate their spending on a single store at the category level, which is consistent with the TSSS view, multiple-store sourcing does not disappear, as would be expected under the TSSS view.[21]

Table 8 presents the results of applying the decomposition separately for each product category. The share of the store-good component decreases for all categories, by about one to two thirds depending on the category. For the majority of categories, the share of the store-basket component increases while the share of the store component decreases.

Table 8: The CCM decomposition by product category

	Carbonated soft drinks	Cereal	Milk	Salty snacks	Yogurt
var(pure store-good)	11%	15%	16%	13%	8%
var(store-basket)	43%	53%	69%	47%	64%
var(store)	56%	41%	33%	49%	31%
2*covar(store-basket, pure store-good)	-7%	-8%	-17%	-6%	-2%
2*covar(store, pure store-good)	2%	3%	2%	3%	0%
2*covar(store, store-basket)	-5%	-4%	-3%	-6%	-2%
sum = var(overall)	100%	100%	100%	100%	100%

Note: All variances and normalized by the total variance, $var(p_{i,t}^{sg} - p_{i,t}^{m})$.

Looking at consumer saving decompositions, the evidence from Table 8 is consistent with the interpretation that consumers take advantage of relative price differences both within and across categories. The pure store-good component explains 24% of the variance in consumer savings across all categories. Turning to the decompositions category by category, this figure falls to 8-16% depending on the category. This decrease in the role of the store-good component is consistent with the TSSS view, but the fact that it is not zero supports

[21]Interestingly, TSSS find a smaller role for multi-store sourcing at the category level: "Across all consumers (whether one- or multi-stop) the share of category spending in the category's second store is 4 percent (panel A3, p.2317)."

the KM/KMRT hypothesis that at least some consumers compare prices of the same good across stores.

5.4 Robustness

We conducted a wide array of robustness tests in order to ensure that our findings are not the result of special circumstances. We provide a summary here. Table 9 replicates the decomposition presented in Table 6 and reports the results of six robustness tests. For each of these tests, Table 10 reports the variances of the KM store-good and the CCM pure store-good components. The 'Baseline' scenario (first column in Table 9 and first row in Table 10) corresponds to the CCM decomposition from Table 6.

Table 9: Robustness check of the CCM decomposition

	Baseline	Filter	Eau	Pitts	SW2	SW3	$var_{i,m,t}(\Delta_{it})$
var(pure store-good)	24%	23%	18%	31%	40%	33%	28%
var(store-basket)	46%	42%	36%	60%	44%	36%	46%
var(store)	49%	52%	61%	31%	49%	49%	50%
2*covar(sb, psg)	-15%	-16%	-12%	-19%	-30%	-15%	-20%
2*covar(s, psg)	4%	4%	7%	0%	32%	3%	3%
2*covar(s, sb)	-7%	-7%	-10%	-3%	-35%	-6%	-8%
sum = var(overall)	100%	100%	100%	100%	100%	100%	100%

Note: *s*, *sb* and *psg* in the covariance names stand for store, store-basket and pure store-good respectively.

Table 10: Variances of the pure store-good and store-good components for each robustness check

	CCM pure store-good	KM store-good
Baseline	24%	54%
Filter	23%	50%
Eau Claire	18%	42%
Pittsfield	31%	72%
Store weight 2	40%	54%
Store weight 3	33%	54%
$var_{i,m,t}(\Delta_{it})$	28%	54%

Note: each row corresponds to a column in Table 9.

The "Filter" column reports results obtained when we filter out panelist-quarter observations with fewer than 20 purchases per quarter. The concern being addressed is that the average purchase count in IRI is smaller than in Nielsen. We want to check that the results do not change when we increase the purchase count per panelist-quarter.

Columns "Eau" and "Pitts" show results for each of the two markets separately. Nielsen contains 54 geographically dispersed markets. One concern is that our two markets may not be representative of the average Nielsen market. Although we are limited in what we can do about this, we can at least check that the results are not driven by a single market. Both markets point to the same conclusion: the KM store-good component is more than twice the size of the CCM pure store-good component.

Columns "SW2" and "SW3" report estimates using alternative store-good weights to the weight $w_{i,j,\cdot,t}$ and $w_{i,\cdot,s,t}$ used in the definition of $p^{sb}_{i,t}$ in Table 2 to weight the store-good prices in the calculation of the store-basket price index. A problem with these weights is that they overestimate the store-basket price index if a good in the panelist's basket has an abnormally high price in a store visited by the panelist. The good may never be bought by the panelist in that store, and for that matter, by most consumers. SW2 assumes that the panelist purchases each good in her basket proportionally to how the average consumer in the market would purchase the good among the stores visited by the panelist. This method takes care of the problem presented above. Another concern is that the panelists' baskets vary from quarter to quarter because the one-quarter window is too short. SW3 computes the weights for the goods in a consumer's basket, $w_{i,j,.,t}$, using a centered three-quarter window.

Finally, the last column considers a different way to aggregate the variance decompositions across markets and quarters. The method reported in the baseline column follows KM's approach: the variance decomposition is conducted by quarter and then aggregated over quarters. The method reported in column $var_{i,m,t}(\Delta_{it})$ computes a single variance decomposition for all panelist-quarter observations.

The results are broadly similar acroos all seven columns in Table 9. In Table 10 we see that in every case, the variance of the CCM pure store-good component is substantially larger than the variance of the KM store-good component, about double the size in several cases. Overall, the large number of robustness tests reported here, along with additional tests reported in the online Appendix, do not produce any evidence against our main conclusions.

5.5 Summary of results

Our key finding is that allowing for the possibility of consumers choosing the best store for the basket shrinks the store-good component to less than half the size found by KM and attributes a larger part of the variance of the HPI to the new store-basket component.

When we conduct the decomposition separately for each product category we find that the store-good component shrinks even more, but does not disappear. This is suggestive of heterogeneous consumer behavior: some consumers shop across stores for individual products, while others make choices at the category level.

Table 13 in Appendix C.2 provides the results of the full CCM decomposition, which includes the transaction component. At 51%, the transaction component from the decomposition of the IRI data is substantially larger than the 16% found by KM with the Nielsen data. This result, which survived a barrage of robustness checks, is an important point of divergence between our results and those of KM. The difference is not due to the methodology, hence it must come down to the use of different data sets, even though it is not clear why consumers in the different panels would behave so differently. We believe that the 51% transaction component that we find is consistent with the recent literature on promotions and its emphasis on heterogeneous consumer behavior (Pesendorfer, 2002; Hendel and Nevo, 2006a,b; Griffith, Leibtag, Leicester, and Nevo, 2009). Reassuringly, the relative magnitudes of the pure store-good, store-basket and store components do not change significantly when the transaction component is included. Nonetheless, the divergence in estimates of the transaction component remains a puzzle for future exploration.

6 Concluding remarks

Price dispersion provides price-conscious consumers with the opportunity to save by shopping around for the best deals. Recent work has documented substantial price dispersion in grocery stores. Combined with the fact that many households spend a significant fraction of their income in grocery stores, this suggests that the scope for savings from grocery shopping is considerable. Consumers can save by searching for the lowest price for identical products both across stores and over time. They can also save by buying in bulk, consuming generic brands or using coupons.

In order to understand the different ways in which consumers save, we adopt and modify the variance decomposition methodology of Aguiar and Hurst (2007) and Kaplan and Menzio (2015). Our modification incorporates the insight that store expensiveness is consumer-specific: one store may be the cheapest place to buy a specific basket of goods, but another store may be the cheapest for a different basket. In practical terms, it amounts to a refinement of the KM decomposition that breaks down their store-good component into two parts that we call pure store-good component and store-basket component. This allows us to address the following question: do (many) consumers really choose the right store for the right product, as KM conclude? Or are they actually just choosing the right store for their basket? The results from our decomposition suggest that the latter is the case. A large fraction of the variance in consumer saving is due to variation in consumers' ability to choose the best store for their basket, and a smaller part is due to variation in ability to choose the right subset of products from the each store. We conclude that the definition of store expensiveness, whether it is consumer specific or common to all consumers, has a significant impact in understanding consumer savings.

Our work adds to a growing literature that attempts to make sense of supermarket pricing and consumer shopping behavior. A branch of this literature mines large store and household datasets to establish stylized facts about pricing and demand. The importance of consumer baskets has long been recognized in the literature on grocery shopping. The latest research – including this paper – is now establishing that consumers vary in their ability to shop for baskets. Conversely, baskets are constrained by the products assortments one can find at the stores one visits. Hitsch, Hortaçsu, and Lin (2021) find that assortments across stores tend to be specialized, and that a similar store assortment within a chain is associated with a similar degree of price dispersion and similar demand elasticities. Understanding how stores tailor product assortments and prices to target specific consumer baskets is an interesting topic that warrants further investigation.

References

Aguiar, M., and E. Hurst, 2007, "Life-Cycle Prices and Production," *American Economic Review*, 97(5), 1533–1559.

Bronnenberg, B. J., M. W. Kruger, and C. F. Mela, 2009, "The IRI Marketing Data Set," *Marketing Science*, 27(4), 745–748.

Ching, A. T., and M. Osborne, 2020, "Identification and Estimation of Forward-Looking Behavior: The Case of Consumer Stockpiling," *Marketing Science*, 39(4), 707–726.

Clerides, S., and P. Courty, 2017, "Sales, Quantity Surcharge, and Consumer Inattention," *The Review of Economics and Statistics*, 99(2), 357–370.

DellaVigna, S., and M. Gentzkow, 2019, "Uniform Pricing in U.S. Retail Chains," *The Quarterly Journal of Economics*, 134(4), 2011–2084.

Dubois, P., and H. Perrone, 2019, "Price Dispersion and Informational Frictions: Evidence from Supermarket Purchases," unpublished manuscript.

Erdem, T., S. Imai, and M. Keane, 2003, "Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1(1), 1–64.

Griffith, R., E. Leibtag, A. Leicester, and A. Nevo, 2009, "Consumer Shopping Behavior: How Much Do Consumers Save?," *Journal of Economic Perspectives*, 23(2), 99–120.

Hendel, I., and A. Nevo, 2006a, "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74(6), 1637–1673.

——— , 2006b, "Sales and Consumer Inventory," *RAND Journal of Economics*, 37(3), 543–561.

——— , 2013, "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 103(7), 2722–2751.

Hitsch, G. J., A. Hortaçsu, and X. Lin, 2021, "Prices and Promotions in U.S. Retail Markets," *Quantitative Marketing and Economics.*

Hosken, D., and D. Reiffen, 2004, "Patterns of Retail Price Variation," *The RAND Journal of Economics*, 35(1), 128–146.

Kaplan, G., and G. Menzio, 2015, "The Morphology of Price Dispersion," *International Economic Review*, 56(4), 1165–1206.

Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter, 2019, "Relative Price Dispersion: Evidence and Theory," *American Economic Journal: Microeconomics*, 11(3), 68–124.

Lach, S., 2002, "Existence and Persistence of Price Dispersion: An Empirical Analysis," *The Review of Economics and Statistics*, 84(3), 433–444.

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

Moen, E. R., F. Wulfsberg, and Ø. Aas, 2020, "Price Dispersion and the Role of Stores," *The Scandinavian Journal of Economics*, 122(3), 1181–1206.

Neslin, S. A., 2002, "Sales Promotion," in *Handbook of Marketing*, ed. by B. Weitz, and R. Wensley. SAGE Publications, pp. 310–338.

Pesendorfer, M., 2002, "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75(1), 33–66.

Pires, T., 2016, “Costly Search and Consideration Sets in Storable Goods Markets,” *Quantitative Marketing and Economics*, 14(3), 157–193.

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

Seiler, S., 2013, “The Impact of Search Costs on Consumer Behavior: A Dynamic Approach,” *Quantitative Marketing and Economics*, 11, 155–230.

Sorensen, A. T., 2000, “Equilibrium Price Dispersion in Retail Markets for Prescription Drugs,” *Journal of Political Economy*, 108(4), 833–850.

Thomassen, Ø., H. Smith, S. Seiler, and P. Schiraldi, 2017, “Multi-category Competition and Market Power: A Model of Supermarket Pricing,” *American Economic Review*, 107(8), 2308–51.

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

Appendices

A Computations of consumer store baskets weights

To compute the store-basket price index, we first compute a store-specific household price index, $p_{i,s}$ (see Panel C in Table 2). To deal with the case where a product purchased by consumer i, is not available in another store visited by consumer i, we assume that the consumer purchases the goods in her basket that are available each store visited, in the same relative proportions as the basket's proportions (see Table 11).[22]

Table 11: Store-basket weights with partial product availability

J_i	Consumer basket	$J_i = \{j \ s.t. \ \omega_{i,j,s} > 0 \ for \ some \ s\}$
S_i	Stores visited	$S_i = \{s \ s.t. \ \omega_{i,j,s} > 0 \ for \ some \ j\}$
$J_{i,s}$	Basket availability	$J_{i,s} = \{j \ s.t. \ \omega_{i,j,s} > 0 \ for \ s \in S_i\}$
$\omega_{i,j,\cdot\vert s}$	Store basket weights	$\omega_{i,j,\cdot\vert s} = \frac{\omega_{i,j,\cdot}}{\sum_{j \in J_{i,s}} \omega_{i,j,\cdot}}$
$p_{i,s}$	Store-specific basket	$p_{i,s} = \sum_{j \in J_{i,s}} \mu_{j,s} \omega_{i,j,\cdot\vert s}$

B Proof of proposition 1

Proposition 1. *A consumer has zero store-good savings, $p_i^{sb} = p_i^{sg}$, when: (a) she visits a single store; (b) store-good prices $\mu_{j,s}$ do not vary across stores visited; or (c) she purchases the same share of each good in all stores visited ($\frac{\omega_{i,j,s}}{\sum_j \omega_{i,j,s}}$ constant across s).*

Proof. Combining the definitions of p_i^{sg} and p_i^{sb} from Table 2, we obtain

$$p_i^{sg} - p_i^{sb} = \sum_{j,s} \mu_{j,s} \left(\omega_{i,j,s} - \omega_{i,\cdot,s} \omega_{i,j,\cdot} \right).$$

To prove claim (a), denote by s_i the single store visited by consumer i. We have $\omega_{i,j,s_i} = \omega_{i,j,\cdot}$, $\omega_{i,j,s} = 0$ for $s \neq s_i$, and $\omega_{i,\cdot,s_i} = 1$. We obtain $p_i^{sg} - p_i^{sb} = \sum_j \mu_{j,s_i} \left(\omega_{i,j,s_i} - \omega_{i,\cdot,s_i} \omega_{i,j,\cdot} \right) = \sum_j \mu_{j,s_i} \left(\omega_{i,j,\cdot} - \omega_{i,j,\cdot} \right) = 0$.

[22] As a technical point, the consumer may not purchase goods in the same proportion in the store-basket and store-good indexes.

Condition (b) says that $\mu_{j,s} = \mu_j$ for all (j, s). We obtain $p_i^{sg} - p_i^{sb} = \sum_j \mu_j \left(\sum_s (\omega_{i,j,s} - \omega_{i,\cdot,s}\omega_{i,j,\cdot}) \right) = \sum_j \mu_j (\omega_{i,j,\cdot} - \omega_{i,j,\cdot}) = 0$.

To prove claim (c) note that condition $\frac{\omega_{i,j,s}}{\sum_j \omega_{i,j,s}}$ constant across s is equivalent to $\omega_{i,j,s} = \omega_{i,\cdot,s}\omega_{i,j,\cdot}$. We conclude that $p_i^{sg} = p_i^{sb} = \sum_{j,s} \mu_{j,s} (\omega_{i,j,s} - \omega_{i,\cdot,s}\omega_{i,j,\cdot}) = 0$. □

Example presented after the proof: To begin, we present a simple class of store pricing policies such that normalized prices take only two values. Consider a market where: (a) all stores pay the same cost for each product, possibly varying from product to product, and then each store chooses a markup for each product that may be low or high, and (b) stores sell the same quantity share of low and high products. Since there is no temporal variation, we can omit without loss of generality the i sub-index on the transaction price $P_{j,s}$, and we also have $\mu_{i,j,s} = \mu_{j,s}$.

Statement (a) says that stores charge price $P_{j,s} = c_j\alpha_e$ for expensive products and $P_{j,s} = c_j\alpha_c$ for cheap ones, where c_j is the cost of product j and $\alpha_e > \alpha_c$ are the markups. Denote by E_j the set of stores where product j is expensive. Applying KM2 (see Panel A in Table 2), we have $P_j = c_j(\alpha_e x_e^j + \alpha_c(1 - x_e^j))$, where $x_e^j = \frac{\sum_{i,s \in E_j} q_{i,j,s}}{\sum_{i,s} q_{i,j,s}}$ is the quantity share of product j sold at an expensive price. According to statement (b) x_e^j is constant across j, $x_e^j = x_e$. Applying KM1, we obtain that the normalized prices are $\mu_{j,s} = \frac{\alpha_e}{x_e\alpha_e + (1-x_e)\alpha_c} \equiv e$ for expensive products and $\mu_{j,s} = \frac{\alpha_c}{x_e\alpha_e + (1-x_e)\alpha_c} \equiv c$ for cheap ones. This establishes the statement that normalized prices can take only two values.

Next, we establish that $\mu_s = 1$ for all s when the expenditure share of expensive products is constant across stores. Applying KM5 from Table 2, we have

$$\mu_s = e\frac{\sum_{i,j \text{ s.t. } s \in E_j} P_{j,s} q_{i,j}}{\sum_{i,j,s} P_{j,s} q_{i,j,s}} + c\frac{\sum_{i,j \text{ s.t. } s \notin E_j} P_{j,s} q_{i,j,s}}{\sum_{i,j} P_{j,s} q_{i,j,s}}$$

which does not depend on s because the expenditure share of expensive products, $\frac{\sum_{i,j,s \in E_j} P_{j,s} q_{i,j,s}}{\sum_{i,j} P_{j,s} q_{i,j,s}}$, is constant across stores. Moreover, plugging the above formula for μ_s in the weighted average $\sum_s \frac{\sum_{i,j} P_{j,s} q_{i,j,s}}{\sum_{i,j,s} P_{j,s} q_{i,j,s}} \mu_s$, we obtain that $\sum_s \frac{\sum_{i,j} P_{j,s} q_{i,j,s}}{\sum_{i,j,s} P_{j,s} q_{i,j,s}} \mu_s = 1$ and conclude that $\mu_s = 1$.

Finally, the store-specific basket price index is $p_{i,s} = \sum_j \mu_{j,s}\omega_{i,j,\cdot}$ and by assumption $\omega_{i,j,\cdot} = \frac{1}{n}$ for n products and $\omega_{i,j,\cdot} = 0$ for the remaining ones. Under the first scenario discussed in the text, we have $\mu_{j,s} = c$ for a single purchased product j and $\mu_{j,s} = e$ for the remaining $n - 1$ products in the consumer's basket. The store-specific basket is $p_{i,s} = \frac{1}{n}c + \frac{n-1}{n}e$ for each store and this is also the value of p_i^{sb}.

C Replication of KM decompositions

We replicate (using the IRI dataset) the KM decompositions for the transaction prices (KM equation 7 in Section 3) and for the household price indexes (KM equation 14 in Section 4, corresponding to equation 1 using the CCM's notations).

C.1 Replication of KM price decomposition (Section 3)

Table 12 reports the results of the price decomposition using the IRI dataset (column one), and for comparison purpose, copies the results of KM price decomposition using Nielsen data (the third column corresponds to Table 3 Column 3 in KM p.14). Columns 2 and 4 renormalize the variances and covariance after ignoring the transaction component.

Table 12: Decomposition at transaction level

	IRI		Nielsen	
	With tran	Without tran	With tran	Without tran
Transaction	65%	–	62%	–
Store-good	31%	89%	30%	81%
Store	4%	11%	7%	19%
2cov(tran,sg)	0%	–	0%	–
2cov(tran,s)	0%	–	0%	–
2cov(sg,s)	0%	1%	0%	0%

The price decompositions are fairly similar across the two datasets. The share of the transaction component is large in both datasets (65% in IRI versus 62% in Nielsen), and similar in both IRI markets (63% in Eau Claire and 66% in Pittsfield), suggesting that promotions play a similar role for our five products categories as it does for the much wider set of products included in KM's analysis.

C.2 Replication of KM household decomposition (section 4)

Table 13 replicates the KM decomposition with the transaction component using the data from IRI. Column one presents the result for the decomposition with the transaction and store-basket components (a combination of equations 1 and 3). Column two re-normalizes the components to obtain the KM decomposition (equations 1). For comparison purpose, Column three copies the values of these components using the data from Nielsen (see KM Column 3, Table 7 in KM p.25). The main difference between the KM decomposition applied to the two different datasets is a significantly higher transaction component in the IRI dataset (51% instead of 16%).

Table 14 shows that the large transaction component is present in both markets and in all five categories. The first column copies column two from Table 13 as a baseline. Column two and and three report the decomposition for the two markets separately. The next five columns replicate the baseline column for the five product categories (carbonated soft drinks, cold cereal, milk, salty snacks and yogurt). The transaction component has the same magnitude in all columns. The same holds if we filter out panelist-quarter observations with fewer than 20 purchases per quarter.

It is difficult to explain why the temporal component explains a larger share of consumer saving in the IRI dataset. Section C.1 has shown that the temporal component explained the same share of price variation in the two dataset. It is not the case that households can take advantage of greater temporal variations (e.g. more frequent or deeper promotions) for the set of products selected from IRI dataset. One explanation could be that households represented in the IRI dataset are more heterogeneous in their ability to take advantage of promotions.

Table 13: Decomposition with transaction component

<table>
<tr><th></th><th>Var-Cov</th><th>Components</th><th>KM-IRI</th><th>KM-Nielsen</th></tr>
<tr><td>$var(\Delta_{i,t}^{a,sg})$</td><td>51%</td><td>Transaction</td><td>51%</td><td>16%</td></tr>
<tr><td>$var(\Delta_{i,t}^{sg,sb})$</td><td>10%</td><td rowspan="2">Store-good</td><td rowspan="2">22%</td><td rowspan="2">53%</td></tr>
<tr><td>$var(\Delta_{i,t}^{sb,s})$</td><td>19%</td></tr>
<tr><td>$var(\Delta_{i,t}^{s,m})$</td><td>20%</td><td>Store</td><td>20%</td><td>39%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{a,sg}, \Delta_{i,t}^{sg,sb})$</td><td>8%</td><td rowspan="2">2cov(tran,sg)</td><td rowspan="2">9%</td><td rowspan="2">5%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{a,sg}, \Delta_{i,t}^{sb,s})$</td><td>1%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{a,sg}, \Delta_{i,t}^{s,m})$</td><td>0%</td><td>2cov(tran,s)</td><td>0%</td><td>1%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{sb,s}, \Delta_{i,t}^{sg,sb})$</td><td>-6%</td><td rowspan="3">2cov(s,sg)</td><td rowspan="3">-1%</td><td rowspan="3">-13%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{s,m}, \Delta_{i,t}^{sg,sb})$</td><td>2%</td></tr>
<tr><td>$2cov(\Delta_{i,t}^{s,m}, \Delta_{i,t}^{sb,s})$</td><td>-3%</td></tr>
<tr><td>Sum=$var(p_{i,t}^{sg} - p_{i,t}^{m})$</td><td>100%</td><td></td><td>100%</td><td>100%</td></tr>
</table>

Table 14: Decomposition with transaction component (robustness)

	KM-IRI	Eau	Pitts	cate1	cate2	cate3	cate4	cate5
Transaction	51%	45%	57%	52%	54%	60%	57%	63%
Store-good	22%	19%	25%	21%	23%	26%	22%	22%
Store	20%	29%	11%	25%	16%	13%	19%	10%
2cov(tran,sg)	9%	9%	9%	3%	7%	2%	3%	6%
2cov(tran,s)	0%	0%	-1%	0%	0%	0%	0%	0%
2cov(sg,s)	-1%	-1%	-1%	-1%	0%	-1%	-1%	-1%