Shopping Frictions with Household Heterogeneity: Theory & Empirics[*]

Krzysztof Pytka[†]

March 2024

Abstract

This paper examines the impact of price dispersion on household consumption, highlighting the role of economic status in shaping purchasing behaviors. Leveraging detailed scanner data, I document high-earning employees pay 1.5 to 7% more than lower-earning ones for the same or similar goods. A causal link between income and the prices is established using the Economic Stimulus Act of 2008. The findings indicate that 8 to 22% of the increase in household spending following a transitory income shock is due to higher prices paid. Despite a broader variety in the consumption baskets of wealthier households, very few goods are tailored to specific income groups. Integrating consumer search with the savings problem, I propose a new model to reconcile the observed patterns and quantify the impact of retail-market frictions on consumption. Counterfactual analysis shows that over two-thirds of households face higher prices due to a price externality.

Keywords: price dispersion, household heterogeneity, NielsenIQ, search externality, data sparsity, fiscal stimulus

JEL codes: E00, E21, L11

[*] I am grateful for useful insights from Árpád Ábrahám, Mark Aguiar, Christian Bayer, Antonio Ciccone, Piero Gottardi, Jeremy Greenwood, Anne Hannusch, Jessie Handbury, Kenneth Judd, Greg Kaplan, Philipp Kircher, Rasmus Lentz, Guido Menzio, Kurt Mitman, Luigi Paciello, Daniel Runge and participants of numerous conferences. Lion Szlagowski's stellar research assistance was invaluable to this work. The previous version of the article was distributed under the title *Shopping Effort in Self-Insurance Economies*. All errors are mine. The researcher's own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†] University of Mannheim; `me@kpytka.eu`

I. Introduction

The traditional income-fluctuation problem, which lies at the heart of workhorse macroeconomic models such as real-business cycle or standard incomplete market models, assume that the law of one price always holds. This means that each product is characterized by a single and unique price. While this assumption lends some tractability to these models, increasing empirical evidence suggests a systematic price dispersion exists. In this paper, I confront this view both empirically and theoretically. To this end, I address the following questions: (*i*) How do the prices paid for the same goods differ across the income distribution, and are consumption baskets systematically different between households?; (*ii*) Which theoretical frameworks are consistent with the observed differences?; (*iii*) What are the macroeconomic implications of these new theories compared to standard consumption models?

This article contributes in two significant ways. On the empirical side, I utilize detailed price scanner data and time-use surveys to document significant heterogeneity in prices for the same or very similar goods, as well as variations in shopping behavior across the income spectrum. A causal relationship between income and prices is established by exploiting the differential timing of the 2008 tax rebates in the US. The observed price adjustments are important for understanding the response of consumption expenditures to liquidity injections. Moreover, consumption baskets between rich and poor consumers are more similar than is typically assumed. Although wealthy households exhibit a higher variety in their consumption baskets, from an aggregate consumption perspective, very few goods are exclusively tailored to a specific income group of consumers. I demonstrate that the documented facts cannot be fully explained by either the precautionary saving model or existing consumer search theories. Motivated by these findings, I propose a novel and tractable framework that integrates random search for consumption within the optimal savings problem, representing the theoretical contribution of this article. This model predicts an equilibrium price distribution characterized by a search externality, where the shopping decisions of one type of household influence the market conditions faced by others, as all households purchase goods from the same markets. The calibrated version of the model is employed to demonstrate the extent of these market externalities.

To investigate price heterogeneity across income brackets, I utilize the Kilts-NielsenIQ Consumer Panel, which tracks 40,000-60,000 American households, capturing detailed scanner data. My analysis spans from 2004 to 2014, covering 630 million transactions for nearly

2 million unique, barcode-level products across 87 million shopping trips. To standardize consumption bundles for comparison, I adopt the methodology outlined by Aguiar and Hurst (2007), constructing an individual price index for each household. This index is calculated as the ratio of actual expenditures to the cost of purchasing the same bundle at average prices paid by other consumers.

With the household price indices at hand, I document how the paid prices vary across different households. Employees with earnings above the median level pay from 1.5 to 7.1% higher prices than those with below-median earnings. The size of the documented differential depends on the definition of the goods. Lower values are associated with more restrictive goods definitions, which, as I show, may be underestimated due to data sparsity. In addition to this, in my analysis employment and retirement are not associated with much lower household price indices. In both cases the price differentials are small and do not exceed 1.4% for any of the considered specifications. This suggests that the differences previously documented in the literature (*cf.,* Kaplan and Menzio, 2015; Aguiar and Hurst, 2007) might be driven by the high-income workers who were pooled with the low-income workers as one homogeneous reference group.

While systematic price heterogeneity across individuals has been documented in numerous studies, to the best of my knowledge, there is no evidence for a causal relationship between economic status and price indices. For this reason, I exploit the random timing of tax rebates from the Economic Stimulus Act of 2008. As I demonstrate in the months following the receipt of a stimulus payment, household price indices increase by between 0.4% and 1.3% compared to the pre-treatment months. This documented increase in the prices paid accounts for between 8% and 22% of the households' rise in overall consumption expenditures. These results shed new light on understanding consumption responses. Typically, the empirical consumption literature obscures the distinction between consumption and expenditures. My findings indicate that price adjustments are a significant driver of individual household consumption responses.

Next, I aim to dissect the differences in the composition of consumption baskets between high-income and low-income workers. Understanding these differences is crucial for constructing an empirically relevant model of the relationships between consumers and retailers. Differences may stem from the *intensive margin*, where wealthier consumers increase the quantity of goods they already purchase, or from the *extensive margin*, which involves diversification in consumption patterns through the acquisition of new types of goods. I demonstrate that *over* 100% of the differences in consumption baskets between poor and rich workers can be explained by the extensive margins. The baskets of rich households exhibit a higher variety, but, when a certain good is consumed, rich households consume *fewer* units of purchased goods. This raises a natural question: with richer households having a wider variety of baskets, is there product specialization towards certain groups of consumers? The evidence suggests very little to no empirical relevance of such specialization. Goods popular among poor consumers are also popular among rich households.

Building upon this analysis, I then turn to examine differences in time spent shopping across households. By leveraging data from the 2003–2018 waves of the American Time-Use Survey and a complementary Well-Being Module, it becomes evident that high-income households allocate approximately 7% more time to shopping activities. This observation, alongside the tendency of wealthier households to incur higher prices, challenges the typical narrative from consumer search literature that associates increased shopping time with lower prices. One potential explanation for this is that wealthier individuals may view shopping as a form of leisure rather than a necessity-driven task, which could account for their tendency to spend more time and pay higher prices. However, this hypothesis finds no empirical support; insights from the Well-Being Module show that shopping-related well-being does not significantly differ between affluent households and other groups.

Motivated by the new findings, I propose a model of consumer search featuring heterogeneous households. This framework incorporates random price search, in the spirit of Burdett and Judd (1983) and Butters (1977), into a standard incomplete-market model as developed by Huggett (1993) and Aiyagari (1994). In this model, all households face idiosyncratic income risk and engage in consumption-savings decisions. Consumption is modeled at the extensive margin, with all goods being perfectly substitutable. Households derive higher utility from the increased variety in their consumption baskets. As consumption expands, households randomly incorporate products not previously consumed within a given period. Each additional item in the basket generates disutility. Shopping is modeled through random price search, with each household endogenously determining their unit price-search intensity and the number of consumed products. Retailers set their prices driven by two motives: capturing surplus from consumers and business stealing from competitors. The equilibrium price distribution emerges as the outcome of a game between retailers and households.

In the calibrated version of the model I show that the price channel accounts for around 8% of the overall consumption responses to transitory shocks, which matches the empirically documented lower bound. Finally, the model economy is used to evaluate the magnitude of price externalities across different households. The externality arises due to the fact that the retailers are not able to distinguish captive customers and bargain hunters. Consequently, search strategies of one individual affects the price distribution of other customers. In a counterfactual analysis, I show that over two thirds of all households pay higher prices due to a negative externality generated by shoppers with low search intensity.

II. Related literature

This paper contributes by proposing a novel and tractable theory of consumer search that aligns with new empirical findings. It connects with several strands of literature.

Price dispersion. Price dispersion has a long tradition in economics, tracing back to a seminal paper by Stigler (1961). However, the empirical macroeconomic literature started exploring this issue more recently. The availability of detailed microeconomic transactional scanner data has enabled researchers to document that retirees and the unemployed pay lower prices than the employed for the same products (Aguiar and Hurst, 2007; Kaplan and Menzio, 2015). My analysis introduces a new dimension of price differences across the income distribution, showing that higher-income individuals tend to pay more than lower-income ones for the same or similar products. Furthermore, utilizing the quasi-experimental context of the Economic Stimulus Act of 2008, I establish a causal relationship between a household's economic status and the prices they pay.

Consumption responses to income shocks. The response of household consumption expenditures to transitory shocks has garnered significant attention from both academics and policymakers. While some studies identify transitory shocks through semi-structural estimations, as done by Blundell et al. (2008) and Commault (2022), a substantial body of literature employs fiscal stimuli to explore the effects of transitory shocks on consumption (Johnson et al., 2006; Parker et al., 2013; Broda and Parker, 2014; Michelacci et al., 2021). Unlike these articles, which consider consumption expenditures (pc) jointly, I distinguish between the impact of transitory shocks on prices (p) and consumption (c).

Heterogeneity in consumption baskets. Previous research has explored price heterogeneity along the income dimension. Majority of observed differences in prices between high-income and low-income households have been attributed to the purchase of different products at higher prices rather than to differences caused by search frictions (Broda et al., 2009; Handbury, 2021; Argente and Lee, 2020). In this article, I observe that a portion of the price differences previously linked to quality can be explained by search frictions, particularly when the issue of data sparsity is directly addressed.

Consumer search. On the theoretical side, this paper introduces a new model of consumer search integrated into a standard income-fluctuation problem. While previous studies have incorporated directed consumer search into consumption-savings problems, such as Bai et al. (2024) or Huo and Ríos-Rull (2015), this paper models purchasing as a consequence of random search. This approach is motivated by the observation that consumers navigate product markets differently, leading to pronounced search externalities. The model adopts the random search framework in the spirit of Burdett and Judd (1983) and more recently Kaplan and Menzio (2016). Unlike these prior works, in this model, households can determine both their consumption level and price search intensity as endogenous decisions and are not assumed to be hand-to-mouth.

The setup of the earlier version of this paper, (Pytka, 2018), has already been used and was extended in various dimensions by other researchers (*cf.* Kang, 2018; Nord, 2023).[1] Nord (2023) adapts the proposed framework to multiproduct markets where goods are tailored to specific consumer segments, with households consuming *all* varieties but in different proportions. Motivated by their empirical observation of limited basket overlap among consumer groups, the extension offers an intriguing yet empirically problematic perspective. While their application of price-search models offers a compelling and thought-provoking approach to reconcile conflicting evidence on the cyclicality of retail markups, my analysis suggests that the structural assumptions of their extension might not withstand scrutiny due to a severe finite-sample bias in their empirical findings.

III. Empirics of Household Shopping

In this section, I analyze the shopping behavior of American households from three different perspectives: paid prices, composition of consumption baskets, and time spent shopping. The primary focus is on the differences among working individuals. The theoretical implications of these findings are discussed jointly in the conclusions of this section.

[1] While the fundamental structure of the economy outlined in the previous version remains largely unchanged, I am now motivated by new empirical findings to take a clearer stance on the interpretation of the consumption basket in the model.

Data. For the analysis of prices and composition of consumption bundles, I utilize the Kilts-NielsenIQ Consumer Panel (KNCP) dataset. The KNCP tracks approximately 40,000 to 60,000 American households, with about 40,000 for the years 2004-2006 and increasing to 60,000 from 2007 onwards. Panelists use in-home scanners or mobile apps to provide NielsenIQ with information about their grocery purchases from any outlet across all US markets. Every purchase of each product is linked to a specific shopping trip made by the household. Additionally, respondents report their socio-demographic characteristics on an annual basis. NielsenIQ provides weights for each household to ensure the sample is representative for the US economy. The sample of households is drawn from 54 geographic markets, known as Scantrack markets. In this analysis, I include data from all markets from 2004 through 2014. During this period, the KNCP collected information on 630 million transactions for nearly 2 million unique products defined at the barcode level, which were purchased in 87 million shopping trips. The analysis of shopping effort utilizes data from the 2003–2018 waves of the American Time Use Survey (ATUS). The ATUS, conducted by the U.S. Census Bureau, randomly selects individuals from a subset of households from the Current Population Survey. Each wave is based on 24-hour time diaries where respondents report their activities from the previous day in specific time intervals. The ATUS staff then categorizes these activities into one of over 400 types.

A. Consumer Prices

Analyzing price differentials among consumers is inherently complex due to the diverse range of goods in their baskets. Variability in both the types and quantities of items complicates direct price comparisons. To accurately document and analyze these differences, I adopt a methodology proposed by Aguiar and Hurst (2007). This technique involves constructing a unique price index for each consumer unit for every period, reflecting the specific quantities and prices of goods they acquire. Essentially, this index is derived by comparing the actual expenditure on a basket to a hypothetical cost, calculated using average prices paid by others for the same (or very similar) items. Next, equipped with this, I explore variations of price indices across different households.

Household price indices (Aguiar and Hurst, 2007). The used methodology follows closely the one proposed by Aguiar and Hurst (2007) with small adjustments. In the original paper the authors focus on households from Denver from January 1993 through March 1995, while in my analysis I study households from all 54 Scantrack markets which are projected on the representative sample of the US population with the use of weights. Consequently, the methodology has been adapted to new features of more recent releases of the KNCP. Products $i \in I$ are bought by households $j \in J$ on shopping trip (date) t in period m. Then the consumption expenditures of households j in month m is given by:

$$X_m^j = \sum_{i \in I, t \in m} p_{i,t}^j q_{i,t}^j. \tag{1}$$

where $p_{i,t}^j$ and $q_{i,t}^j$ are paid prices and quantities of product i, respectively. The average price of product i in period m in market r weighted by the number of purchases is as follows:

$$\bar{p}_{i,m}^r = \sum_{j \in J(r), t \in m} w_{j,m} p_{i,t}^j \left(\frac{q_{i,t}^j}{\bar{q}_{i,m}} \right), \tag{2}$$

where $\bar{q}_{i,m}$ is the overall number of purchases, $\bar{q}_{i,m} := \sum_{j \in J, t \in m} w_{j,m} q_{i,t}^j$ made by households (denoted by $J(r)$) from market r. The computed statistics are meant to be representative thanks to weights, $w_{j,m}$, provided by NielsenIQ, which sum up to the total number of households in the US.[2] The hypothetical cost of consumption of the household j from market $r(j)$ if she paid the average prices for purchased products would be given by:

$$Q_m^j = \sum_{i \in I, t \in m} \bar{p}_{i,t}^{r(j)} q_{i,t}^j. \tag{3}$$

Then the household price index $\bar{P}_{j,t}$ can be obtained from:

$$\bar{P}_{j,m} := \frac{\frac{X_m^j}{Q_m^j}}{\sum_{j' \in J(r)} w_{j',m} \frac{X_m^{j'}}{Q_m^{j'}}}. \tag{4}$$

What is a Good? The value of the household index, $\bar{P}_{j,m}$, is influenced by the definition of each good, as reflected through the average prices of goods purchased by other households, $\bar{p}_{i,m}^{r(j)}$. Thus, the employed empirical strategy requires making a stand about the definition

[2]For the sake of clarity and to make the exposition of the presented formulas easier to follow, I introduce a standardization, denoted as $w_{j,m} = \frac{\text{Projection Factor}_{j,m}}{\sum_{j \in J(r)} \text{Projection Factor}_{j,m}}$.

of a good in terms of its physical characteristics (i), the market in which it was bought ($r(j)$), and the timing of the transaction (m). As I argue, those choices are quite important. On one hand, overly broad definitions of goods may aggregate products that should not be considered close substitutes. On the other hand, excessively narrow definitions inevitably bias the household indices toward a value of one.[3] The latter issue can be particularly significant in the used dataset.

Regarding physical characteristics (i), are a 20-oz. bottle of Coke and a 12-oz. can of Coke the same product or different ones? Are a 12-oz. can of Coke and a 12-oz. can of Pepsi the same product or different ones? Some products exhibiting similar features may be very close substitutes to each other and it might be recommended to study them jointly. In order to be as agnostic as possible about what a good is I consider two conservative definitions. First, I define products at the bar-code level. In the second scenario, products featuring very similar characteristics are pooled together as one product.[4] Some of goods are described by 19 characteristics, such as flavor, package type, size, organic claims, or amount of salt. Goods sharing *exactly all* characteristics the same are considered as one product. Products not described by NielsenIQ are still studied at the bar-code level. After pooling goods with similar characteristics, the number of unique products drops from initial 1,990,173 to 473,879.

Another choice that must be made in defining a good is the market region (r). Should the price of 12-oz. can of Coke purchased in New York be compared with the price of 12-oz. can of Coke purchased in Scranton? I consider two extreme cases, the market is defined either at the Scantrack level (*e.g.*, New York Coke is not comparable with Coke purchased in Scranton) or nationwide (*e.g.*, New York Coke is compared with Coke bought in Los Angeles).

Unpleasant data sparsity. Table 1 presents a share in the aggregate consumption of purchased goods with different numbers of observed transactions depending on the definition of the goods.[5] As observed, under the narrowest definition of a good (barcode, Scantrack

[3]By construction, if good i is purchased only once and only by one household j, the average price equals the price from the sole transaction in trip t within the designated market $r(j)$, *i.e.*, $\bar{p}^{r(j)}_{i,m} = p^j_{i,t}$.

[4]In this scenario a 12-oz. can of Coke and a 12-oz. can of Pepsi are the same products as they have the same package, the same volume, and the same flavor, while 20-oz. bottle of Coke and a 12-oz. can of Coke are different products.

[5]For the product definitions, I chose to group transactions by month to mitigate unnecessary seasonal price fluctuations, which can occur even within quarterly periods. For instance, turkey sold just before Thanksgiving in November and turkey sold in December should not be considered close substitutes.

Table 1: Number of transactions for different goods and shares in the aggregate consumption. The table reports the ratios of implied aggregate consumption after excluding products that were purchased less than 2, 10, and 20 times in a given month to the implied aggregate consumption without any restrictions.

Product definition (i)	Market definition (r)	Period (m)		No. of transactions		
			≥ 1	≥ 2	≥ 10	≥ 20
Similar features	Nationwide	Monthly	1	0.990	0.942	0.903
Similar features	Scantrack market	Monthly	1	0.860	0.517	0.404
Bar code	Nationwide	Monthly	1	0.963	0.814	0.723
Bar code	Scantrack market	Monthly	1	0.713	0.287	0.209

market), transactions of products bought only once represent almost 30% of the aggregate consumption. Meanwhile, less than 30% of the aggregate consumption involves products purchased 10 times or more. As previously mentioned, a positive share of transactions for products purchased only once can bias household price indices toward one. Therefore, while more inclusive definitions of goods might raise questions on the level of substitutability of different goods, the narrow definitions might underestimate the level of heterogeneity. Nonetheless, even the most inclusive definition of goods applied in this study is quite conservative. All in all, the results are reported for four combinations of definitions of goods and markets, with a remark that estimates for more restrictive definitions are lower bounds of the true price heterogeneity across households.[6] Figure 1 depicts the distribution of price indices according to those definitions. The discussed data sparsity poses additional challenges in deconstructing the heterogeneity in consumption bundles. As demonstrated in Subsection B of Section III, the structure of the dataset could result in an overestimation of consumption polarization (or underestimation of basket overlap) if the analysis is not approached with the necessary caution.

A.1. High Earners Pay Higher Prices

I explore variations in prices paid across different households by regressing the logarithms of household price indices, $\ln \bar{P}_{j,m}$, on dummies identifying various household groups. These include households with total annual income above the median, non-employment status for

[6]There exists an alternative approach to address the issue of products with few transactions, which centers on focusing only on products purchased with sufficient frequency. However, as demonstrated in Appendix A, this method significantly compromises the representativeness of the dataset for capturing the dynamics of aggregate consumption.

Figure 1: Distribution of household price indices. The figure displays histograms of individual price indices, $\bar{P}_{j,m}$, as defined in Equation (4), for various goods definitions. Histograms for more restrictive definitions show greater concentration around the value of one, partly due to a higher occurrence of rare transactions, as detailed in Table 1.

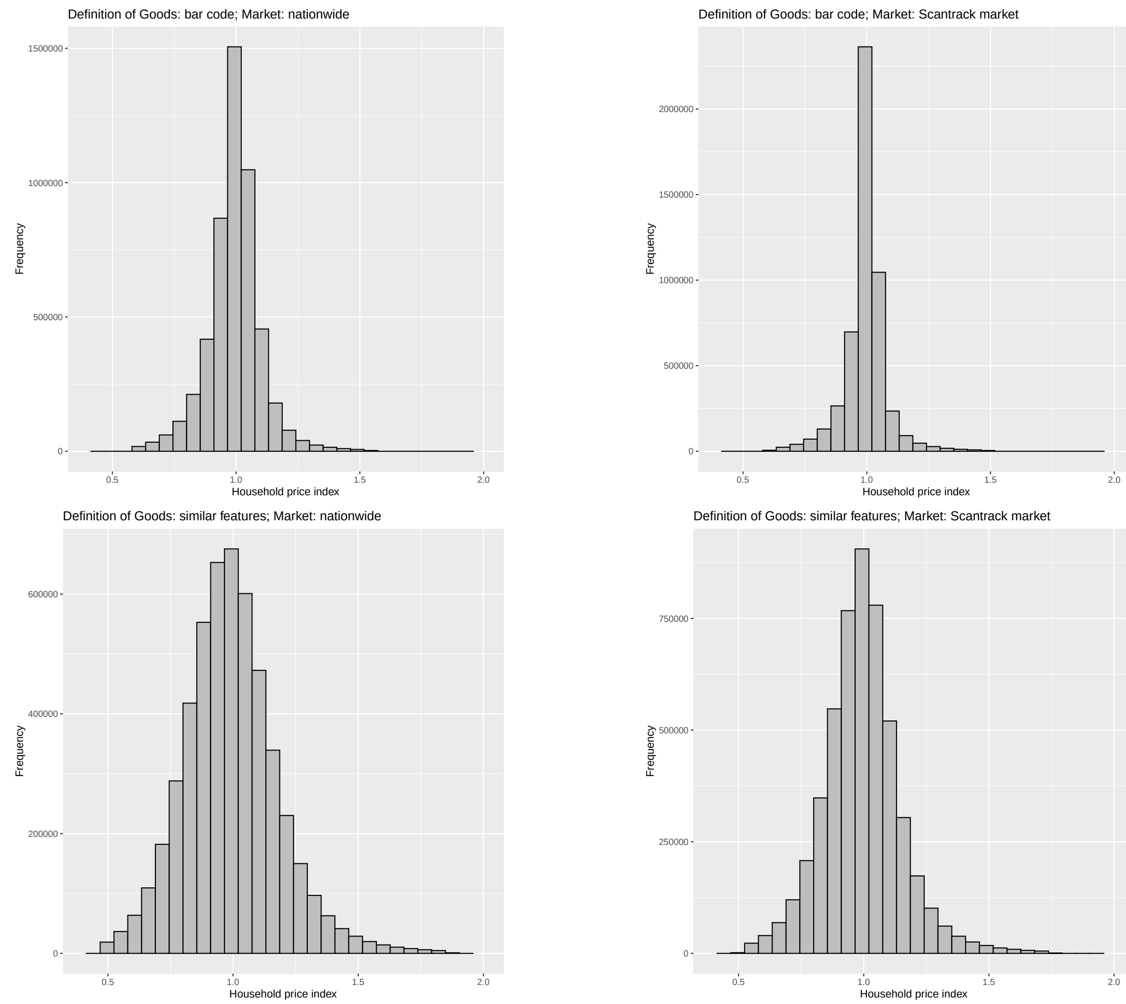

either head of household of working age (two dummies), retirement status for either head of household of working age (two dummies), household composition (eight dummies), age dummies (for each head), year and month dummies, and Scantrack region dummies. The primary explanatory variable is the one accounting for households with total annual income above the median. In this specification, the reference group comprises households with total annual income at or below the median. Non-employment and retirement status dummies are included for comparative purposes. In the KNCP, the income variable is reported for 2 years prior to the panel year; therefore, I utilize future reports from households 2 years ahead to obtain the current income. Households not present in the panel in future years are excluded.[7] Table 2 presents estimates of the regression for all four specifications of the household price indices. There are two immediate striking results.

Firstly, households with higher earnings pay between 1.5% and 7.1% more than those with lower earnings. These price differentials are comparable to or exceed those documented for non-employed versus employed households (between 0.8% and 4.6%, as documented by Kaplan and Menzio (2015)) and for retired versus working-age households (3.6% at the barcode level, as reported by Aguiar and Hurst (2007)). All estimates are highly statistically significant. Unlike previous studies, my analysis uses employed households earning below the median as the reference group. In this specification, the price differential for retired and non-employed groups diminishes considerably. The impact of non-employment or retirement on prices does not exceed 1.5% compared to prices paid by lower earners. This finding suggests that the income distribution among employed consumers is at least as crucial as the extensive margins on the labor market in understanding variations in consumer prices.

What is even more important, the conditional price heterogeneity across different households increases when we consider less restrictive (but still quite conservative) definitions of goods.[8] In previous studies, some price heterogeneity across the income distribution was also observed (Broda et al., 2009; Handbury, 2021). Handbury (2021) documented that products (defined at the barcode level) consumed by households in the highest income bracket (annual income higher than $100,000) are around 5% more expensive compared to the same products purchased by households with an annual income below $25,000. Furthermore, this price gradient increased to 17% when products were defined at more general module levels, such as processed cheese slices American, carbonated soft drinks, etc. The overall conclusion, similar to Broda et al. (2009), was that most of the price differential is explained by the fact that affluent consumers buy *different* and more expensive products rather than the same goods at higher prices. In the employed methodology, the differential is identi-

[7] Such restrictions may introduce selection issues, potentially compromising the sample's representativeness. However, as demonstrated in Appendix A, these concerns do not affect the dataset's overall representativeness.

[8] Just to illustrate that feature aggregation remains highly conservative, consider the definition of goods in this dataset: regular 12-oz processed American cheese sold in foil wraps of brand A and regular 12-oz processed American cheese sold as slices of brand A are still classified as two distinct products. All characteristics, including flavor and USDA certifications, are exactly the same; the only difference is in the packaging — in the former case, slices are wrapped in foil, whereas in the latter, slices are separated by pieces of paper. The only substitute for the former is regular 12-oz processed American cheese sold in foil wraps of brand B.

fied using transactions for products that are frequently purchased by different households. However, the used dataset exhibits massive sparsity if the products are defined at the most granular level, which may lead to an underestimation of the magnitude of price heterogeneity. As I stated before, grouping similar products might mitigate this issue. The findings presented in Table 2 reveals that a portion of the previously observed price heterogeneity among households can be attributed to high-income households potentially paying more for identical or extremely similar products, as opposed to selecting different and systematically more expensive products.

Table 2: Household price indices across different income and employment states

	$\ln \bar{P}_{j,m}$ (1)	(2)	(3)	(4)
HH Earnings > median(HH Earnings)	0.020***	0.015***	0.071***	0.052***
	(0.002)	(0.002)	(0.002)	(0.002)
Non-employed in working age (Male)	−0.007***	−0.006***	−0.014***	−0.010***
	(0.001)	(0.001)	(0.002)	(0.002)
Non-employed in working age (Female)	−0.007***	−0.004***	−0.010***	−0.006***
	(0.001)	(0.001)	(0.002)	(0.001)
Retired (Male)	−0.002	0.0001	−0.00002	−0.001
	(0.002)	(0.002)	(0.004)	(0.003)
Retired (Female)	0.002	0.001	0.002	0.001
	(0.002)	(0.002)	(0.004)	(0.003)
HH composition dummies	*Yes*	*Yes*	*Yes*	*Yes*
Age dummies (both heads)	*Yes*	*Yes*	*Yes*	*Yes*
Month dummies	*Yes*	*Yes*	*Yes*	*Yes*
Year dummies	*Yes*	*Yes*	*Yes*	*Yes*
Scantrack market dummies	*Yes*	*Yes*	*Yes*	*Yes*
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	5,084,254	5,084,254	5,084,254	5,084,254
Number of panelists	150,153	150,153	150,153	150,153
R^2	0.034	0.016	0.071	0.042

Notes: Robust standard errors clustered at the household and year level are included in parentheses.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

The relationship between economic status and paid prices was further explored using information on financial liquidity reported by households in the 2008 wave. Analogous to the findings presented in this section, poorer households tend to pay lower prices on average. Discussion of these findings is relegated to Appendix B. Additionally, shopping amenities could potentially give rise to price differences between poor and rich households.

To quantify the impact of shopping amenities on price disparities, I created expenditure-weighted household indices to assess the expensiveness level of the stores visited. This methodology is an extension of that proposed by Kaplan and Menzio (2015). The analysis indicates that the store-specific component contributes minimally, accounting for only 17 to 30% of the overall heterogeneity in household price indices between the affluent and the less affluent. Details of the analysis of shopping amenities can be found in Appendix C.

A.2. Prices Are Causally Related to Income

While there is quite rich and robust evidence on systematic heterogeneity in price indices across different households, in fact, we do not know much about the causal nature of those differences. To my knowledge, no previous studies have thoroughly addressed this issue. In particular, one could not rule out a possibility that some other confounding factors can drive both paid price and economic status of the households. This section aims to: (i) establish a causal link between the price indices of individual consumption bundles and the income levels of those households; and (ii) assess the impact of households' price adjustments on their expenditure responses to Economic Stimulus Payments (ESP).

Research Design. I leverage the quasi-experimental framework provided by the Economic Stimulus Act of 2008, which involved distributing tax rebates to approximately 130 million eligible taxpayers. Eligible households received payments as tax rebates, with amounts ranging from $300 to $600 for single filers and $600 to $1,200 for married couples filing jointly.[9] Due to the scale of the whole program, the ESPs could not be paid at once. For this reason, some randomization in the timing of disbursement had to be introduced. This randomization was achieved by linking the disbursement week to the last two digits of recipients' Social Security numbers, effectively rendering the timing of payments a random assignment. As such, the ESPs can be interpreted as exogenous income shocks, offering valuable insights into the causal effects of income on individual price indices. For this analysis, I utilize data from a 2008 tax rebate survey conducted by NielsenIQ on behalf of Broda and Parker (2014), merging it with information from the KNCP used in previous subsections.

[9] A detailed discussion on the program's structure is available in Sahm et al. (2010) and Parker et al. (2013).

To assess the pass-through of income shocks on the prices paid by households, I estimate the following regression:

$$\ln \bar{P}_{j,m} = \alpha_j + \beta_{-1} \cdot \sum_{s=-1}^{-3} R_{j,m-s} + \beta_0 \cdot \sum_{s=0}^{2} R_{j,m-s} + \beta_1 \cdot \sum_{s=3}^{5} R_{j,m-s} + \beta_2 \cdot \sum_{s=6}^{8} R_{j,m-s} + \eta_m + \varepsilon_{j,m}, \quad (5)$$

where the dependent variable, $\ln \bar{P}_{j,m}$, represents the log price index for household j in month m. The term α_j denotes the household fixed effect, while η_m represents the fixed effect of each month. The key independent variable, $R_{j,m}$, is a dummy variable indicating receipt of the payment in month m. The coefficient β_0 corresponds to the average price response to the ESP during the month of receipt and the two subsequent months. Coefficient β_{-1} measures the average price response in anticipation of the ESP up to three months before receipt. Coefficients β_1 and β_2 capture the average price response 3-5 months and 6-8 months after receipt, respectively. In this specification, periods at least four months before receipt of the ESP serve as the reference group.

Table 3 presents the results from estimating Equation (5) across four definitions of goods previously discussed. The estimates compellingly indicate that income level causally affects the prices paid, with its impact enduring up to 8 months post-receipt of the tax rebate. Notably, the ESP, averaging \$900 in the dataset, led to a price increase for identical or very similar goods by between 0.4 and 1.3%. Furthermore, the statistically insignificant β_{-1} suggests that anticipation of the ESP did not influence current consumption.

While the presented estimates might seem modest at first glance, it is essential to recognize that the ESP shocks were relatively minor in comparison to overall household income. Furthermore, the estimated price response reflects changes in prices across the *entire* consumption baskets. To fully understand the magnitude of the identified impact, examining the contribution of this price response to the overall response of household consumption expenditures is instructive. Let us assume a household receives the ESP in month τ. The expected response of expenditures in month $\tau + s$ to the tax rebate can then be broken down into two components:

$$\mathbb{E} \ln \left(\frac{\bar{P}_{j,\tau+s} Q_{j,\tau+s}}{\bar{P}_{j,\tau-1} Q_{j,\tau-1}} \right) = \underbrace{\mathbb{E} \left(\ln \bar{P}_{j,\tau+s} - \ln \bar{P}_{j,\tau-1} \right)}_{\text{Price channel}} + \underbrace{\mathbb{E} \left(\ln Q_{j,\tau+s} - \ln Q_{j,\tau-1} \right)}_{\text{Consumption channel}}, \quad (6)$$

Table 3: Price response to the ESP

Response to the ESP	$\ln \bar{P}_{j,m}$			
	(1)	(2)	(3)	(4)
Quarter before, β_{-1}	0.002	0.001	0.002	0.003
	(0.001)	(0.001)	(0.002)	(0.002)
Quarter of receipt, β_0	0.006***	0.004*	0.009**	0.008***
	(0.002)	(0.002)	(0.003)	(0.003)
One quarter after, β_1	0.008***	0.005**	0.009**	0.011***
	(0.003)	(0.002)	(0.005)	(0.004)
Two quarters after, β_2	0.008**	0.006**	0.011**	0.013***
	(0.003)	(0.003)	(0.006)	(0.005)
Month dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	345,768	345,768	345,768	345,768
Number of panelists	29,289	29,289	29,289	29,289
R^2	0.589	0.549	0.605	0.565

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

where $Q_{j,m}$ represents the quantity of the composite good as defined by Aguiar and Hurst (2007), previously introduced in Equation (3). The first component of the decomposition, representing the price channel of the response, $\mathbb{E}\left(\ln \bar{P}_{j,\tau+s} - \ln \bar{P}_{j,\tau-1}\right)$, equates to $\beta_i - \beta_{-1}$ from Equation (5), with $i \in \{0, 1, 2\}$ corresponding to the quarter associated with month s post-ESP receipt. To ascertain the price channel's contribution to the overall response, it is essential to also consider the remaining elements, *i.e.*, $\mathbb{E} \ln \left(\frac{\bar{P}_{j,\tau+s} Q_{j,\tau+s}}{\bar{P}_{j,\tau-1} Q_{j,\tau-1}}\right)$ and $\mathbb{E}\left(\ln Q_{j,\tau+s} - \ln Q_{j,\tau-1}\right)$. Consequently, I adapt the model specified in Equation (5) to include alternative dependent variables, substituting $\ln \bar{P}_{j,m}$ with $\ln \left(\bar{P}_{j,m} Q_{j,m}\right)$ and $\ln Q_{j,m}$.[10]

Table 4 shows the decomposition of the expenditure responses into two components, price changes and consumption changes. As it turns out, households' adjustments in the relative paid prices account for between 8 and 22% of the total changes in the overall consumption expenditures after receiving the tax payment.

[10]The proposed decomposition effectively breaks down the growth in expenditures, and thanks to the properties of logarithms, the covariate terms are canceled out. An alternative decomposition for $\mathbb{E}[\bar{P}_{j,\tau+s} Q_{j,\tau+s} - \bar{P}_{j,\tau-1} Q_{j,\tau-1}]$ is presented, with the main conclusions remaining unchanged. The derivation of this alternative decomposition, along with all regression tables necessary for both decompositions, is included in Appendix D.

Table 4: Decomposition of the expenditure responses to the ESP

Product aggregation	Area aggregation	Price channel: $\frac{\mathbb{E}\left(\ln \bar{P}_{j,\tau+s} - \ln \bar{P}_{j,\tau-1}\right)}{\mathbb{E}\ln\left(\frac{\bar{P}_{j,\tau+s}Q_{j,\tau+s}}{\bar{P}_{j,\tau-1}Q_{j,\tau-1}}\right)}$			Consumption channel: $\frac{\mathbb{E}(\ln Q_{j,\tau+s} - \ln Q_{j,\tau-1})}{\mathbb{E}\ln\left(\frac{\bar{P}_{j,\tau+s}Q_{j,\tau+s}}{\bar{P}_{j,\tau-1}Q_{j,\tau-1}}\right)}$		
		QTR_0	QTR_1	QTR_2	QTR_0	QTR_1	QTR_2
Bar code	Nationwide	12.5%	11.6%	12.0%	87.5%	88.4%	88.0%
Bar code	Scantrack	8.1%	8.5%	10.0%	91.9%	91.5%	90.0%
Features	Nationwide	22.2%	15.3%	18.1%	77.8%	84.7%	81.9%
Features	Scantrack	16.8%	16.3%	19.0%	83.2%	83.7%	81.0%

A.3. Price Variance in a Single Transaction is Higher for High Earners

In addition to studying variation of price indices across households, I look into the variance of relative prices $\frac{p_{i,m}}{\bar{p}_{i,m}}$ in a *single* transaction for both high earners and low earners. As shown in Table 5, a single transaction of high earning households is subject to higher risk for all definitions of a good. While this result may look rather technical, as I discuss later, it provides a desired property of a micro-founded search protocol for theoretical models. This result is discussed jointly with other findings at the end of the whole section.

Table 5: Price variance of a single transaction for high earners and low earners.

Product aggregation	Area aggregation	$\text{Var}\left(\frac{p^r_{i,m}}{\bar{p}^r_{i,m}} \mid j \text{ is High Earner}\right)$	$\text{Var}\left(\frac{p^r_{i,m}}{\bar{p}^r_{i,m}} \mid j \text{ is Low Earner}\right)$
Similar features	Nationwide	0.508	0.407
Similar features	Scantrack market	0.396	0.320
Bar code	Nationwide	0.178	0.152
Bar code	Scantrack market	0.156	0.128
	Number of transactions	207 millions	258 millions
	Number of customers	45,901	60,646

B. *Deconstructing Heterogeneity in Consumption Bundles*

In this section, I aim to dissect the differences in the composition of consumption baskets between high-income and low-income workers. I argue that a thorough understanding of these differences is crucial for constructing an empirically relevant model of the relationships between consumers and retailers. On the one hand, differences may stem from the *intensive margin*, where wealthier consumers increase the quantity of goods they already purchase. On the other hand, heterogeneity may arise from the *extensive margin*, which entails a diversification in consumption patterns through the acquisition of new types of goods. Furthermore, it is also important to ascertain whether there are specialized goods that are systematically consumed only by certain groups of consumers, or whether all consumers purchase similar goods, leading to their shopping actions being pooled together in the same retail market.

To this end, I propose a formal decomposition into intensive and extensive margins. The average consumption (in units) of product i by household j belonging to group k, denoted as $\mathbb{E}(c_i^k)$, can be expressed as $\mathbb{E}(c_i^k) = \Pr(c_i^{j\in k} > 0) \cdot \mathbb{E}(c_i^k | c_i^k > 0, j \in k)$, where $\Pr(c_i^{j\in k} > 0)$ represents the probability that a household from group k consumes product i, and $\mathbb{E}(c_i^k | c_i^k > 0, j \in k)$ is the conditional expected consumption of this product, given that it is purchased.[11] Given this, the difference in the consumption of good i between high-earnings households and low-earnings households, $\mathbb{E}(c_i^{\text{Rich}}) - \mathbb{E}(c_i^{\text{Poor}})$, can be decomposed in the following way:

$$\mathbb{E}(c_i^{\text{Rich}}) - \mathbb{E}(c_i^{\text{Poor}}) = \underbrace{(\Pr(c_i^{\text{Rich}} > 0) - \Pr(c_i^{\text{Poor}} > 0)) \cdot \mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Rich})}_{\text{Extensive Margin}_i} + \underbrace{(\mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Rich}) - \mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Poor})) \cdot \Pr(c_i^{\text{Poor}} > 0)}_{\text{Intensive Margin}_i}. \tag{7}$$

Figure 2a shows the expenditure-weighted distribution of the extensive margin, defined as $\frac{\text{Extensive Margin}_i}{\mathbb{E}(c_i^{\text{Rich}}) - \mathbb{E}(c_i^{\text{Poor}})}$ from Equation (7). Let $\mathbb{E}_i[\cdot]$ denote the expenditure-weighted average. On average, the contribution of the extensive margin accounts for 113% of the difference between the consumption patterns of rich and poor workers, expressed as $\mathbb{E}_i[\mathbb{E}(c_i^{\text{Rich}}) - \mathbb{E}(c_i^{\text{Poor}})]$. While the contribution of the extensive margin exceeding unity might seem counterintuitive at first, its interpretation is quite straightforward. Specifically, rich households consume a broader variety of products, resulting in a higher average probability of consumption for every good, *i.e.*, $\mathbb{E}_i[\Pr(c_i^{\text{Rich}} > 0)] > \mathbb{E}_i[\Pr(c_i^{\text{Poor}} > 0)]$. Conversely, the number of units consumed—conditioned upon consumption of the good—is *lower* for rich households than for poor workers, namely, $\mathbb{E}_i[\mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Rich})] < \mathbb{E}_i[\mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Poor})]$. Figure 2a also

[11] This formula is a direct consequence of applying the law of iterated expectations, *i.e.*

$$\mathbb{E}(c_i^k) = \Pr(c_i^{j\in k} > 0) \cdot \mathbb{E}(c_i^k | c_i^k > 0, j \in k) + \Pr(c_i^{j\in k} = 0) \cdot \mathbb{E}(c_i^k | c_i^k = 0, j \in k),$$

where the latter term disappears because $\mathbb{E}(c_i^k | c_i^k = 0, j \in k)$ is always equal to 0. In this subsection, products are defined on a nationwide basis at the barcode level, and the time frequency is set to annual; this ensures that my results are directly comparable to those reported by Nord (2023).

illustrates heterogeneity in the contribution of the extensive margin. Notably, for 78% of the aggregate consumption, the contribution of the extensive margin is greater than or equal to 78%. This indicates that the extensive margin is the predominant factor driving the differences in consumption baskets between rich and poor workers.[12]

A natural question that arises from the observation that richer households have a higher variety of baskets is whether there is product specialization towards some groups of consumers. Do we observe a quality ladder, where richer households consume increasingly customized goods which replace the more generic ones purchased by poorer consumers, or rather, are more popular goods universally preferred? Figure 2b compares the probabilities that a certain product is consumed by high-earning consumers versus low-earning consumers. As can be seen, there is an extremely strong positive relationship between these probabilities. The expenditure-weighted correlation is equal to 0.997 and the expenditure-weighted cosine similarity is equal to 0.952.[13] This implies that in the considered dataset, there is no such thing as product specialization towards certain income groups.[14] Products more popular among poor customers are also more popular among rich customers.[15]

[12]Naturally, there is an alternative decomposition to Equation (7), where the extensive margin is defined by $(\Pr(c_i^{\text{Rich}} > 0) - \Pr(c_i^{\text{Poor}} > 0)) \cdot \mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Poor})$, and the intensive margin is given by $(\mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Rich}) - \mathbb{E}(c_i^k | c_i^k > 0, j \in \text{Poor})) \cdot \Pr(c_i^{\text{Rich}} > 0)$. For this decomposition, the contribution of the extensive margin is still greater than unity, being equal to 1.01.

[13]Weighted cosine similarity (in the Euclidean space) of consumption probabilities between high-income, $\Pr(c_i^{\text{Rich}} > 0)$, and low-income workers, $\Pr(c_i^{\text{Poor}} > 0)$, is defined as:

$$\frac{\sum_{i=1}^{n} Agg\_Exp_i \cdot \Pr(c_i^{\text{Rich}} > 0) \cdot \Pr(c_i^{\text{Poor}} > 0)}{\sqrt{\sum_{i=1}^{n} Agg\_Exp_i \cdot \Pr(c_i^{\text{Rich}} > 0)^2} \cdot \sqrt{\sum_{i=1}^{n} Agg\_Exp_i \cdot \Pr(c_i^{\text{Poor}} > 0)^2}}, \tag{8}$$

where Agg_Exp_i are projection-weighted aggregate expenditures on product i of both household groups. Cosine similarity is widely used for comparing multidimensional vectors (*e.g.,* Jaffe, 1986; Hwang et al., 2010; Hristakeva, 2022; Gentzkow et al., 2019a; Chen et al., 2022).

[14]This result should not be interpreted as denying the existence of product quality or a quality ladder at all; it simply indicates that these aspects are not observed along the income distribution. Specialization along other dimensions remains a possibility. For instance, consumers' historical exposure to certain product groups might shape their preferences, leading them to develop a taste for more specialized goods. However, such considerations are beyond this paper's main focus, which primarily examines differences in consumption patterns across income groups.

[15]In a companion paper (Pytka ⓡ Runge, 2024), we explore this phenomenon by formally assessing polarization. We estimate separate consumption models for households in the top and bottom quintiles of the income distribution using penalized multinomial models, which are designed to address high-dimensional choices amidst severe data sparsity. These estimated consumption generating processes are then compared using polarization measures proposed by Gentzkow, Shapiro, and Taddy (2019b). Our results closely mirror those presented in this paper, indicating an absence of polarization.

Figure 2: Heterogeneity in consumption and comparison between rich and poor households.

(a) Heterogeneity in consumption bundles: extensive vs. intensive margins.

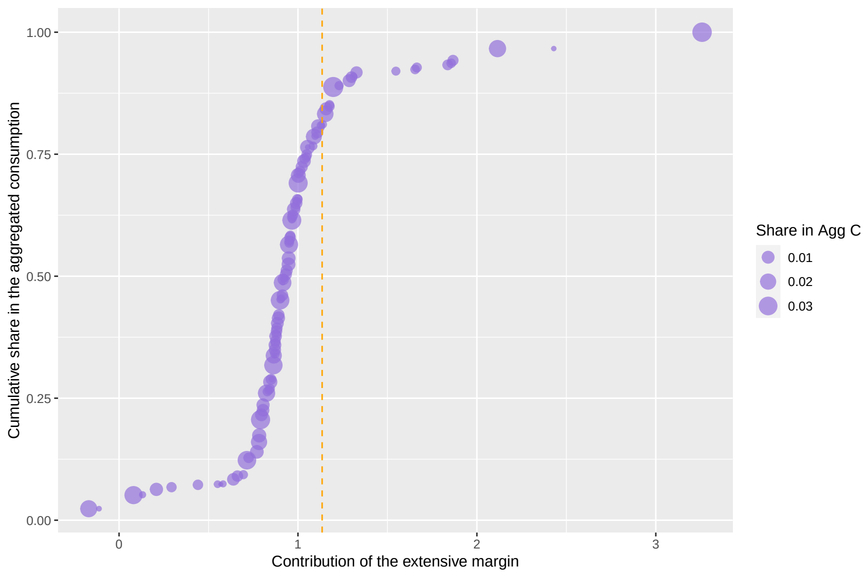

(b) Consumption Comparison of Different Goods: Rich vs. Poor Households

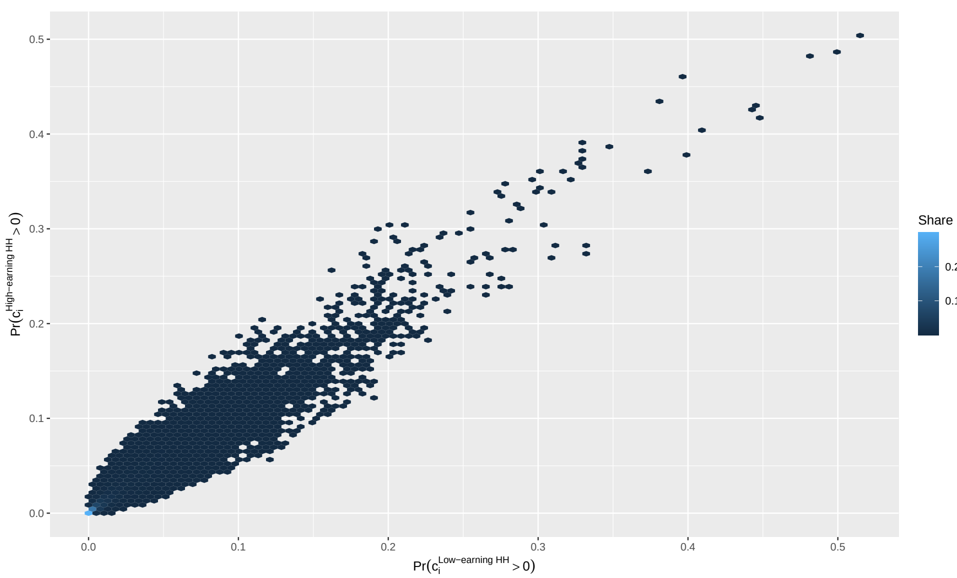

Top panel: The contribution of the extensive margin is calculated as the ratio of the extensive-margin component from Equation (7) to $\mathbb{E}(c_i^{\text{Rich}}) - \mathbb{E}(c_i^{\text{Poor}})$ for each product. Products are arranged on the x-axis according to their extensive margin contributions, from lowest to highest. Purple bubbles represent module averages, with bubble size indicating the share in aggregate consumption. The y-axis shows the cumulative share of aggregate consumed goods where the extensive margin is smaller than the corresponding x-axis value. An orange dashed line marks the overall mean value across all products, identified as 1.1359. Bottom panel: This part of the figure contrasts the proportions of low-income households consuming a specific good against those of high-income households consuming the same item. It is notable that a significant probability mass centers around goods that are purchased very infrequently. Despite this, a positive relationship exists between those variables, even after accounting for this by using expenditure-weighted correlation or cosine similarity from Equation (8).

The presented findings diverge from the analysis independently conducted on the same dataset by Nord (2023). In Appendix E, I explore the nuances of their basket-overlap measure, especially in scenarios characterized by transaction sparsity of products, previously discussed and detailed in Table 1. A permutation test indicates that their estimated measure of basket overlap understates the actual overlap in these contexts and almost three-quarter of the documented effect is due to the finite-sample bias. In Section F of the appendix, the same test applied to my decomposition from Equation (7) reveals minimal to no bias.

C. Shopping Effort: High Earners Spend More Time Shopping

The final aspect of heterogeneity in shopping decisions that I intend to investigate pertains to shopping effort. Specifically, I am interested in whether households paying different prices also demonstrate varying levels of shopping effort. Following established literature (e.g., Aguiar and Hurst, 2007; Aguiar et al., 2013; Kaplan and Menzio, 2016), I adopt time spent shopping as a proxy for shopping effort. To examine differences in shopping time across households, I utilize the American Time Use Survey (ATUS) dataset. My analysis involves analyzing time diaries that detail how American households allocate their time. My focus particularly lies on understanding the relationship between shopping effort and labor market status, including unemployment, retirement, and labor earnings levels. I find that, among employed households, there is a positive correlation between the level of shopping effort and their earnings.

I investigate variations in shopping time across households by using regression analysis on the shopping time against variables such as total annual income above the median, non-employment, retirement status, year and age group dummies, and specific 'shopping needs.' Following Aguiar and Hurst (2007), 'shopping needs' are adjusted for family composition differences by including dummies for partnership status, partner's employment status, and number of children. The analysis considers cumulative daily shopping time (in minutes) for acquiring goods or services (excluding education, restaurant meals, and medical care) and related travel. Activities include grocery shopping, warehouse and mall shopping, banking, haircuts, price/product research, and online shopping. The sample is limited to individuals aged 25 to 75, excluding the self-employed.

Table 6 presents the estimation results. Relative to employed individuals with below-median earnings (the reference group), all other groups allocate more time to shopping, averaging an additional 2-2.5 minutes daily. This increment, statistically significant, represents approximately 7% of the average daily shopping duration of 38 minutes. Estimates for retirees and non-employed working-age individuals are about 7 minutes per day, aligning with findings reported in existing literature.

Table 6: Shopping time across different individuals

	Shopping time		
	(1)	(2)	(3)
Earnings>median(Earnings)	2.590***	2.049***	2.116***
	(0.450)	(0.451)	(0.446)
Nonemployed (in working age)	6.700***	6.712***	6.710***
	(0.508)	(0.508)	(0.503)
Retired	7.916***	7.644***	7.955***
	(0.755)	(0.791)	(0.783)
Age categories	*Yes*	*Yes*	*Yes*
Shopping needs	*No*	*Yes*	*Yes*
Year and day dummies	*No*	*No*	*Yes*
N	149,797	149,797	149,797
R^2	0.010	0.011	0.033

Notes: ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

Previous studies have associated shopping time with price search intensity, indicating that unemployed individuals and retirees spend more time shopping, which, according to consumption search theories, translates to lower prices. However, my findings reveal a seemingly counterintuitive and contradictory observation: households with higher earnings spend more time shopping yet pay lower prices. A potential explanation might be that, for wealthier households, shopping is considered more of a leisure activity, in contrast to other consumers for whom it represents non-market work. Nevertheless, a supplementary study on the well-being implications of shopping, presented in Section G, falsifies this hypothesis.

D. Mapping the Data to (Existing) Models

The empirical findings presented here are interesting in several dimensions. First, they demonstrate that price heterogeneity is substantial, extending beyond the extensive labor margin previously documented in the literature. In fact, the results suggest even greater heterogeneity among workers across the income distribution. Additionally, the observation that higher earners spend more time shopping yet pay higher prices stays at odds with existing random price-search theories, such as those by Burdett and Judd (1983) and more recently Kaplan and Menzio (2016). These theories suggest that greater search efforts should lead to lower prices. This would suggest that theories relying on the directed search might be a better representation of price search. In this class of models, it is possible that consumers with higher earnings and higher consumption decide to choose retailers that have shorter queues (for each unit of consumption) but higher prices. Per unit of consumption those consumers spend less time shopping but overall due to higher level of consumption they may spend more time shopping overall. Nonetheless, those models have some other limitations. The directed search assumes perfect knowledge about prices. Each household directs their shopping activity to stores with different price levels. If this were true then there should not be a systematic difference in the price variance of a single purchase across households with different level of income, which I documented. Moreover, in models of directed search, different consumers visit different markets.[16] This aspect does not align with my findings, which indicate that very few products are tailored to specific income groups and that the level of store expensiveness remains rather similar across different income groups.

Consequently, the empirical analysis suggests that currently there is no one micro-founded representation of shopping that would reconcile all presented findings on price differential across different earnings groups.

IV. In Search of the Theory

Motivated by all findings documented in the empirical section and the absence of a theory that reconciles all observed patterns, I propose a model of consumer search with heterogeneous households. This framework integrates random price search, as described by Burdett and Judd (1983), into the standard incomplete-market model in the tradition of Huggett (1993) and Aiyagari (1994), further enhanced by an overlapping generations structure with life-cycle components, akin to Ríos-Rull (1995). Time is discrete. All households live for a finite number of periods, face idiosyncratic income risk with a deterministic life-cycle component, and make consumption-savings decisions. Consumption is modeled at the extensive

[16] Just to name but a few, Burdett, Shi, and Wright (2001); Acemoglu and Shimer (1999); Menzio and Shi (2011); Bai, Ríos-Rull, and Storesletten (2024).

margin, with perfectly substitutable goods available. Households derive higher utility from the increased variety in their consumption baskets. As households expand their consumption, they randomly add products not previously consumed (in a given period).[17] Including an additional item into the basket requires effort, generating disutility. Shopping is modeled through random price search, with each household endogenously determining their unit price-search intensity and the number of consumed products. A fixed measure of ex-ante identical retailers set their prices based on a combination of two motives: appropriating surplus from customers and stealing business from competitors. The price distribution is an equilibrium outcome of a game between retailers and households.

A. Search for Goods

Within a single period, to purchase a unit of consumption, a household must embark on a shopping trip. A higher number of varieties, denoted as c_t, entails an increased number of shopping trips. The total cost of the bundle, $\mathscr{B}(c_t, s_t)$, is the sum of numerous price lotteries, *i.e.*

$$\mathscr{B}(c_t, s_t) = \int_0^{c_t} p(i)di, \tag{9}$$

where $p(i)$ represents the price obtained during a single shopping trip. A single price is a random variable characterized by the distribution function $p(i) \sim_{i.i.d.} F(p; s_t)$, where s_t is the price search intensity individually chosen by households. Moreover, beyond market acquisitions, goods can also be produced domestically through labor endowment, following a linear home production function: $\zeta \to 1$.

Let $G(p)$ represent the equilibrium (yet to be determined) c.d.f. of prices set by retailers. The price lottery confronting a household is modeled as a compound lottery:

$$F(p; s_t) = (1 - s_t) \underbrace{G(p)}_{\text{Captive purchase}} + s_t \underbrace{\left(1 - [1 - G(p)]^2\right)}_{\text{Non-captive purchase}}. \tag{10}$$

In this model, s_t indicates the likelihood that the household's search for a given variety

[17]The exclusive focus on the extensive margin and the absence of income-group specialization directly stem from findings presented in Subsection B of Section III and reconfirmed in Pytka ⓡ Runge (2024).

results in receiving competing offers from two retailers, with the choice falling on the lower of those offers.[18] Conversely, with the complementary probability, the household is captive, receiving only a single offer to which the consumer is committed.[19]

Even though there is uncertainty in a single transaction and consumers do not know the prices that will be drawn, the fact that this is a repeated activity and consumers make a continuum of trips of measure c_t within each period allows me to deterministically pin down the cost of the consumption basket, $\mathscr{B}(c_t, s_t)$. This is the subject of the following lemma.

Lemma 1 (Cost of consumption bundle)**.** *Let the effective price of a purchase be distributed according to the cdf $F(p; s_t)$. Then the cost of consumption c_t given search intensity converges almost surely:*

$$\int_0^{c_t} p(i)di \xrightarrow{\text{a.s.}} \underbrace{c_t \cdot \mathbb{E}(p|s_t)}_{\mathscr{B}(c_t,s_t)}, \tag{11}$$

where $\mathbb{E}(p|s_t)$ is the average effective price of consumption and is equal to $\mathbb{E}(p|s_t) = \int pdF(p; s_t)$.

Proof. The lemma is an immediate result of applying the weak law of large numbers for random continuum in a version proposed by Uhlig (1996, Theorem 2). □

Next, we are in a position to derive a formula for the average effective price, $\mathbb{E}(p|s_t)$, which emerges as a linear function of the price search intensity, s_t, with its slope and intercept being two sufficient statistics derived from the distribution of prices set by sellers, $G(p)$.

Proposition 1 (Linearity of the Average Effective Price Function)**.** *Given the distribution of quoted prices $G(p)$, the average effective price paid by households is a linear function with respect to the search intensity s:*

$$\mathbb{E}(p|s_t) = p^0 - s_t \cdot MPB, \tag{12}$$

where:

i. $p^0 = \int x\, dG(x)$ represents the price for the fully captive consumer;

[18] Clearly, if p' and p'' are two i.i.d. draws from $G(p)$, then $Pr(x \geq \min\{p', p''\}) = (1 - G(p))^2$. Thus, the cdf of the minimum of two prices is given by $Pr(x \leq \min\{p', p''\}) = 1 - [1 - G(p)]^2$.

[19] Given the focus on non-durable consumption, I assume that households are compelled to buy the good at the offered price, interpreting this as a need for an uninterrupted inflow of products for survival in every "instant" of the period.

ii. $MPB = \mathbb{E}(\max\{p', p''\}) - p^0 (\geq 0)$ *is the marginal (price) benefit of increasing the search intensity* s_t*, where* $\mathbb{E}(\max\{p', p''\})$ *is the expected maximum of two independent draws of prices, equal to* $\int_0^\infty 1 - [G(x)]^2 \, dx$*.*

Proof. The proof is relegated to Appendix H.A. □

As I will demonstrate in subsequent sections, Lemma 1 and Proposition 1 significantly simplify the household's decision-making problem. The total cost of the consumption bundle purchased by each household, denoted as $\mathscr{B}(c_t, s_t)$, is the product of the number of varieties, c_t, and the average effective price, $\mathbb{E}(p|s_t)$, which is a linear function. The intercept of this function corresponds to the expected average price for a household with the probability of two draws being zero and every shopping trip resulting in matching with only one retailer. Households can reduce the average price by increasing their search intensity, s_t. The marginal effect on the average effective price is constant and equals the marginal price benefit, MPB.

Shopping effort vs. monetary cost. The disutility from obtaining goods, $g(c_t, s_t)$, depends on the number of goods, c_t, and the price search intensity, s_t. I assume that $\frac{\partial}{\partial c_t} g(c_t, s_t) > 0$ and $\frac{\partial}{\partial s_t} g(c_t, s_t) > 0$, indicating that both margins increase disutility. Greater price search intensity and a larger number of goods both contribute to increased shopping disutility. Moreover, I assume that $\frac{\partial^2}{\partial c_t \partial s_t} g(c_t, s_t) > 0$. Each unit of consumption brings more disutility if the search intensity is higher, implying that each unit of consumption entails greater disutility if the search intensity is higher.[20] Particularly, I adopt the following functional form for the disutility function:

$$g(c_t, s_t) = \frac{\omega}{1 + \phi} \left(\frac{1 + s_t}{1 - s_t} \cdot c_t \right)^{1+\phi}, \tag{13}$$

where $\phi > 0$ is the price search cost parameter. This disutility function increases with both c_t and s_t, with the cross-partial derivative being positive. Figure 3 visualizes this disutility function across various values of s_t. Within any given period, households have limited resources stemming from their accumulated savings and current income, to be detailed later. Consequently, we observe an inherent tension between shopping effort and monetary cost. The more time and effort households invest in shopping, the more likely they are to find price deals, leading to lower $\mathbb{E}(p|s_t)$, yet incurring a higher disutility, $g(c_t, s_t)$.

[20] This assumption allows me to reconcile the finding that households paying higher prices also spend more time shopping with their purchase of a more diverse variety of items in their basket.

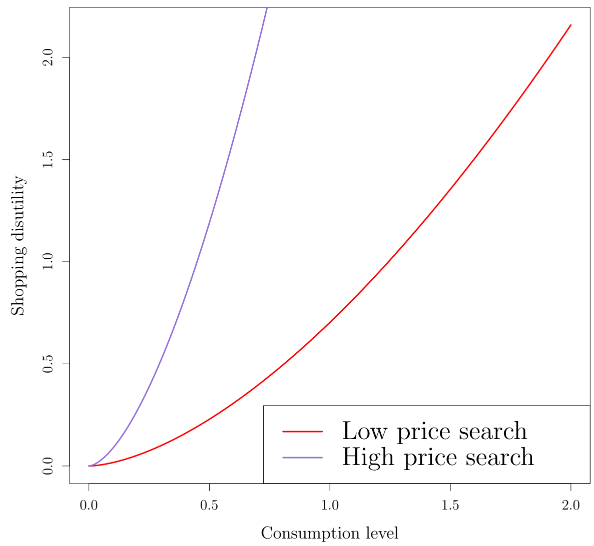

Figure 3: Disutility from shopping as described by the functional specification in Equation 13. The graph illustrates the trade-off faced by households when deciding on the number of goods to purchase (x-axis) against the intensity of price search efforts within the interval [0,1]. The imposed Inada-like condition, $\lim_{s\to 1^-} \frac{\partial}{\partial s} g(c,s) = +\infty$, ensures that a search intensity equaling to one is never optimal. This restriction is crucial for avoiding degenerate Diamond (1971)-type equilibria, as discussed subsequently.

B. Problems of Individual Agents

In the proposed economy, there are two types of agents: a continuum of households and a double continuum of retailers.[21] After deciding how to split their income between savings and expenditures, households engage in purchasing by visiting goods markets subject to search frictions. Retailers, on the other hand, sell goods to the visiting households and set prices based on the distribution of shopping strategies.

Problem of the Households. The model period is one year. The stationary economy is populated by a continuum of households living T periods. Consumers work for T_{work} periods and next go into retirement for $T - T_{work}$ periods. Households have preferences over stochastic sequences of consumption and overall shopping effort $\{c_t, s_t\}_{t=1}^{T}$, represented by

[21] The economy features a continuum of varieties, and for each variety, there is a continuum of sellers, leading to a double continuum of retailers overall. The 'double-continuum' structure in price search models is discussed in more detail by Mangin and Menzio (2024, Footnote 4) or Menzio (2023).

the instantaneous utility function:

$$\underbrace{u(c_t)}_{\text{Consumption}} - \underbrace{g(c_t, s_t)}_{\text{Shopping Effort}}, \tag{14}$$

and the discount factor β. Households aim to maximize expected utility and derive utility from the variety in their consumption baskets, denoted by c_t.[22] In particular, I assume the functional form $\frac{c_t^{1-\sigma}-1}{1-\sigma}$ for $u(c_t)$ and Equation 13 for $g(c_t, s_t)$. While being active in the labor market ($t \in \overline{1, T_{\text{work}}}$), every household faces idiosyncratic wage risk. Log productivities follow an exogenous stochastic process:

$$\ln y_t = \kappa_t + \eta_t + \varepsilon_t, \tag{15}$$

$$\eta_t = \eta_{t-1} + \nu_t,$$

where $\varepsilon_t \sim_{\text{iid}} \mathcal{N}(0, \sigma_\varepsilon^2)$ and $\nu_t \sim_{\text{iid}} \mathcal{N}(0, \sigma_\nu^2)$. The deterministic part, κ_t, represents a lifecycle component common to all households. The martingale part, η_t, and the serially uncorrelated part, ε_t, account for the permanent and transitory components of productivity, respectively. Households older than T_{work} receive a deterministic retirement benefit, which is a function of their last working-age period income with a replacement rate $repl$:

$$\ln y_t = \ln(\text{repl}) \cdot \{\kappa_{T_{\text{work}}} + \eta_{T_{\text{work}}} + \varepsilon_{T_{\text{work}}}\}.$$

Households can hold a single risk-free asset that pays a net return, r. Let a_{t+1} denote the amount of the asset carried over from t to $t+1$. Every household faces a sequence of intertemporal budget constraints:

$$\underbrace{\mathbb{E}(p|s_t) \cdot c_t}_{\mathscr{B}(c_t, s_t)} + a_{t+1} \leq w y_t + (1+r) a_t, \quad \forall_{t \in \overline{1,T}}.$$

Furthermore, each household is subject to an exogenous borrowing constraint $a_{t+1} \geq \underline{B}$.

Having outlined all relevant aspects from the household's perspective, I am now in a po-

[22] More formally, the consumption basket comprises a sum of perfectly substitutable varieties, $\int_0^{c_t} 1 di$. The assumption of substitutability stems from observations in Subsection B of Section III and from polarization analysis conducted in another paper (Pytka ⓡ Runge, 2024). Each household chooses varieties of measure c_t, randomly selected from the set of all available varieties. As households value variety over quantity, each variety can be included in a particular household's consumption basket once or not at all in a given period.

sition to present the dynamic problem faced by a household of age t, characterized by the state $x = (a, \varepsilon, \nu, \eta)$, in the following recursive formulation:

$$\mathcal{V}_t(a, \varepsilon, \eta) = \max_{c,s,a'} u(c) - g(c, s) + \beta \mathbb{E}_{\eta'|\eta} \mathcal{V}_{t+1}(a', \varepsilon', \eta') \tag{16}$$

s.t.

$$\begin{aligned}
\mathscr{B}(c, s) &\leq (1 + r)a + wy - a', \\
\mathscr{B}(c, s) &= (p^0 - s \cdot MPB) \cdot c, \\
a' &\geq \underline{B}, \\
s &\in [0, 1], \\
\log y &= \begin{cases} \kappa_t + \eta + \varepsilon, & \text{for } t \leq T_{work}, \\ log(repl) \cdot \{\kappa_{T_{work}} + \eta_{T_{work}} + \varepsilon_{T_{work}}\}, & \text{for } t > T_{work}, \end{cases} \\
\eta' &= \eta + \nu'.
\end{aligned}$$

Problem (16), at its core, bears resemblance to the textbook consumption-saving problem found in many macroeconomic models. The novel aspect here is that goods' prices are no longer considered exogenous. Households, through their shopping decisions, can influence the prices they pay. Reducing prices is more costly for larger baskets. The relationship between the basket's cost and the disutility from acquiring it is represented by both $\mathscr{B}(c, s)$ and $g(c, s)$. The microfoundations of the search were discussed in the previous subsection. Thanks to Proposition 1, the problem remains relatively tractable and can be solved using standard dynamic programming methods. This allows for the simplification where I do not need to track the entire price distribution; instead, only two distributional statistics, p^0 and MPB, are sufficient.

Retailers' Problem. There is a fixed measure of sellers of mass 1. Each seller is visited θ times.[23] Each seller incurs a unit cost for selling a unit of any product variety. The seller's objective is to set a price p that maximizes profits, $S(p)$:

$$S(p) = \theta \sum_{t=1}^{T} \int \underbrace{\left(1 - \frac{2s_t(x)}{1 + s_t(x)} G(p)\right)}_{\substack{\text{Business} \\ \text{Stealing}}} \underbrace{(p - 1)}_{\substack{\text{Surplus} \\ \text{Appropriation}}} \overbrace{\frac{c_t(x)(1 + s_t(x))}{\sum_{t=1}^{T} \int c_t(x)(1 + s_t(x)) d\mu_t(x)}}^{\substack{\text{Weighting factor} \\ \text{for search distribution, } s_t}} d\mu_t(x), \tag{17}$$

[23] Due to households randomly selecting varieties for their baskets, the problem for each variety is symmetric. Thus, for simplicity, the variety index is omitted in further discussions.

where $\mu_t(x)$ is the distribution of households of age t over the individual states $x = (a, \varepsilon, \nu, \eta)$. In fact, my model extends the retailer's problem described in a classic paper by Burdett and Judd (1983), with an adjustment to account for household heterogeneity. Each retailer maximizing their profit is driven by a tug of war between two competing forces: the desire to charge a higher price to extract more surplus from the consumer and the desire to charge a lower price to steal more consumers from their competitor. The former motive is reflected by the term $(p-1)$, indicating the profit margin over the cost, while the latter is represented by $\left(1 - \frac{2s_t(x)}{1+s_t(x)}G(p)\right)$, denoting the probability that a visiting consumer receives a competing offer priced above p. In equilibrium, these forces find balance across the entire spectrum of the equilibrium price distribution.

C. Equilibrium Characterization

Having outlined the building blocks of the economy, I am now in a position to define the equilibrium of this economy.

Definition 1 (Rational Stationary Equilibrium)**.** *A stationary equilibrium is a sequence of consumption and shopping plans $\{c_t(x), s_t(x)\}_{t=1}^{T}$, and the distribution of quoted prices $G(p)$ and paid prices $F(p; s_t(x))$, distribution of households $\mu_t(x)$ and interest rate r such that:*

1. $c_t(x), s_t(x)$ *are optimal given* r, w, $G(p)$, $\underline{B}$;
2. *individual and aggregate behavior are consistent:*

$$\theta = \sum_{t=1}^{T} \int (1 + s_t(x))c_t(x)d\mu_t(x); \tag{18}$$

3. *retailers post prices* p *to maximize the sales revenues taking as given households' behavior;*
4. *the private savings sum up to an exogenous aggregate level* $\overline{K}$:

$$\sum_{t=1}^{T} \int a_t(x)d\mu_t(x) = \overline{K}; \tag{19}$$

5. $G(p)$ *and* $F(p; s_t(x))$ *are consistent given the household distribution* $\mu_t(x)$;

6. $\mu_t(x)$ is consistent with the consumption and shopping policies.

The concept of equilibrium in the setup considered here does not significantly deviate from the standard one found in the family of models initiated by Huggett (1993) and Aiyagari (1994).[24] Since production is not modeled explicitly, my framework aligns more closely with Huggett's, albeit extended to include an overlapping generations structure similar to Ríos-Rull (1996).[25] A key distinction is that households no longer act as price takers in the goods market. Instead, the price dispersion is an equilibrium object being a result of a game between retailers and households, exactly like in Burdett and Judd (1983).

One major difference from the Burdett and Judd (1983)-type economy, which features up to two dispersed equilibria, is related to consumer search behavior. In their model, consumers are indifferent between drawing only one price quote, facing a higher expected price, and incurring a search cost for a second quote to potentially pay a lower expected price.[26] The dispersed equilibrium arises at the intersection of the search cost and the benefit of conducting a second search.[27] This phenomenon does not occur in my model due to the strictly increasing and convex shopping disutility in relation to the probability of drawing two prices, as dictated by the functional form specified in Equation (13).[28]

The dispersed distribution of posted prices is consistent with the solution to the maximization of the retailers' net sales revenue from Equation (17). Lemma 2 presents properties of an equilibrium of this kind. The proof of the lemma is similar to ones used in Burdett and Judd (1983) and Kaplan and Menzio (2016).

[24] Despite the growing popularity of this class of models, the literature still lacks analytical results. Knowledge about the properties of aggregate consumption and saving behavior is limited, complicating theoretical analysis. I acknowledge that equilibrium multiplicity might potentially emerge, inheriting the structure of aggregate savings from the incomplete-markets model. Only recently have some theoretical results on equilibrium uniqueness under specific conditions been proposed (*e.g.*, Achdou, Han, Lasry, Lions, and Moll, 2021; Light, 2020), yet the question remains open for more general setups.

[25] I depart from the zero aggregate wealth in an endowment economy proposed by Huggett (1993). I do this because the size of aggregate savings directly affects households' consumption decisions. For this reason, I assume that there is an exogenously given supply of assets in the economy, denoted as $\overline{K}$. Next, I ensure that the discount factor, β, is set to match the aggregate wealth-to-income ratio of 2.5, given the observed interest rate, r. Similar deviations from a closed-economy general equilibrium have been made by others, including Kaplan and Violante (2010) and Guerrieri and Lorenzoni (2017).

[26] In their setup, search decision is indivisible. Due to this, ex-ante identical consumers become ex-post heterogeneous. This mechanism is very similar to the indivisible labor supply proposed by Rogerson (1988).

[27] The authors show that the benefit of conducting a second search has a hump-shaped relationship with aggregate search intensity, leading to multiple dispersed equilibria (See Figure 1, Burdett and Judd, 1983).

[28] This property is discussed in more detail in my thesis (Pytka, 2017, Chapter 2).

Lemma 2 (Characterization of the Equilibrium Price Dispersion)**.** *The c.d.f. $G(p)$ exhibits following properties:*

i. $G(p)$ is continuous.

ii. supp $G(p)$ *is a connected set.*

iii. the highest price charged by retailers is equal to ζ,

iv. all retailers yield the same profit, $\forall_{p\in \text{supp } G(p)} S(p) = S^*$,

where supp $G(p)$ *is the smallest closed set whose complement has probability zero.*

Proof. The proof relegated to Appendix H.B. □

For further characterization, it is convenient to decompose the aggregate measure of shopping trips, θ, into two components: the number of trips where customers are captive, $\Psi_{(-)}$, and the number of trips where customers receive competing offers for the same variety, $\Psi_{(+)}$:

$$\Psi_{(-)} := \sum_{t=1}^{T} \int c_t(x)(1 - s_t(x))d\mu_t(x), \tag{20}$$

$$\Psi_{(+)} := \sum_{t=1}^{T} \int c_t(x)2s_t(x)d\mu_t(x), \tag{21}$$

$$\theta = \sum_{t=1}^{T} \int c_t(x)(1 + s_t(x))\mu_t(x) = \Psi_{(-)} + \Psi_{(+)}. \tag{22}$$

Note that $\Psi_{(-)}$, as defined in (20), represents the aggregate measure of visits where customers are captive. In contrast, $\Psi_{(+)}$ in (21) captures instances where households draw two prices and select the lower one. θ, from (22), quantifies the total measure of aggregate shopping as defined in (18), comprising both $\Psi_{(-)}$ and $\Psi_{(+)}$ components. Accordingly, from the retailers' perspective, $\frac{\Psi_{(-)}}{\theta}$ and $\frac{\Psi_{(+)}}{\theta}$ reflect the probabilities of a purchase being captive or accompanied by an alternative offer, respectively. By design, all transactions within $\Psi_{(-)}$ are concluded because buyers have no alternatives in these instances. Conversely, only half of the offers in $\Psi_{(+)}$ lead to sales, as consumers are presented with two price quotes and opt for the more favorable one. Properties from Lemma 2 can be used to derive the formula for an equilibrium price dispersion.

Theorem 1 (Equilibrium Price Dispersion)**.** *Given aggregate statistics of households' shopping decisions* $\{\Psi_{(-)}, \Psi_{(+)}, \theta\}$, *(where* $\Psi_{(-)}, \Psi_{(+)} > 0$*), the equilibrium price dispersion can be expressed in a closed form:*

$$G(p) = \begin{cases} 0, & \text{for } p < \underline{p}, \\ \frac{\theta}{\Psi_{(+)}} - \frac{\Psi_{(-)}}{\Psi_{(+)}} \cdot \frac{\zeta - 1}{p - 1}, & \text{for } p \in [\underline{p}, \zeta], \\ 1, & \text{for } p > \zeta, \end{cases} \tag{23}$$

where the lower bound of supp $G(p)$ *is* $\underline{p} = \frac{\Psi_{(+)}}{\theta} + \frac{\Psi_{(-)}}{\theta}\zeta$.

The closed-form solution for the equilibrium price dispersion, as outlined in Theorem 1, emerges directly from applying the properties of equilibrium described in Lemma 2. This equilibrium is unique in its consistency with equilibrium properties and ensures equal profits across every point of the price support.

There are two layers of information asymmetry from a single retailer's perspective. Firstly, just like in Burdett and Judd (1983), a retailer does not know whether a visiting customer has received a competing offer from another retailer within the same variety market. Moreover, retailers are unaware of a specific visiting customer's probability of drawing two prices.[29] This lack of knowledge stems from the empirical finding that there is no evidence of product specialization towards specific income groups. As a result, retailers are only informed of the average aggregate probability of two draws, $\frac{\Psi_{(+)}}{\theta}$, which suffices to establish the equilibrium price distribution.[30]

Given a price p, the equilibrium c.d.f., $G(p)$, is a linear function that decreases with both the odds of being matched with a captive customer, denoted as $\frac{\Psi_{(-)}}{\Psi_{(+)}}$, and the probability that a visiting buyer receives an alternative offer, represented by $\frac{\Psi_{(+)}}{\theta}$. Consider two economies with identical aggregate shopping efforts θ but different mass of non-captive shopping trips,

[29] All consumers in this model visit the same markets, an important distinction from standard directed search models where different agents visit different markets (Burdett, Shi, and Wright, 2001; Menzio and Shi, 2011; Acemoglu and Shimer, 1999; Moen, 1997).

[30] It is important to note that from Equations (20) - (22), we get the odds: $\frac{\Psi_{(+)}}{\Psi_{(-)}} = \frac{\frac{\Psi_{(+)}}{\theta}}{1 - \frac{\Psi_{(+)}}{\theta}}$.

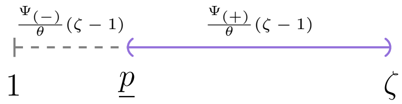

Figure 4: The equilibrium support of $G(p)$ (purple segment). The figure illustrates the lower bound of the equilibrium price dispersion, $\underline{p}$, as a weighted sum of two extreme solutions: the fully competitive (normalized to 1) and the fully monopolistic, which equals ζ. The higher the fraction of non-captive purchases $\frac{\Psi_{(+)}}{\theta}$, the closer the lower bound gets to the competitive solution. Conversely, the higher the fraction of captive purchases, $\frac{\Psi_{(-)}}{\theta}$, the closer $\underline{p}$ is to the monopolistic solution.

$\Psi'_{(+)} > \Psi''_{(+)}$. Notice that:

$$\frac{\partial G(p)}{\partial\left(\frac{\Psi_{(+)}}{\theta}\right)} = \frac{\left(\frac{\zeta-1}{p-1}\right) - 1}{\left(\frac{\Psi_{(+)}}{\theta}\right)^2} > 0$$

for every p within the interior of supp $G(p)$. Thus, the price lottery in the economy with higher search intensity, $\Psi'_{(+)}$, first-order stochastically dominates that in the economy with lower search intensity, $\Psi''_{(+)}$. This observation implies that economies with more intensive search activities exhibit lower expected values of the price lottery. This result aligns with economic intuition: the greater the fraction of buyers receiving competing offers, the stronger the competition among retailers, leading to more aggressive pricing strategies. This result is particularly interesting from a modeling perspective because the aggregate price search intensity, an endogenous equilibrium variable, plays a crucial role in determining the level of markups. Contrary to the standard models of monopolistic competition, such as the one presented by Dixit and Stiglitz (1977), where markups are constant under flexible pricing, this model allows for average prices to fluctuate between competitive and monopolistic extremes. As illustrated in Figure 4, the overall outcome is influenced by the distribution of households' purchasing strategies.

In Figure 5, it is even more evident that the proposed economy exhibits a search externality. The price moments crucial for each household's search problem, namely p^0 and MPB, are derived from the equilibrium price dispersion resulting from the strategic interplay between households and retailers. These moments are influenced by the price c.d.f., $G(p)$, which in turn depends on $\frac{\Psi_{(+)}}{\theta}$ as stated in Equation (23). An individual household has the flexibility to choose any average price for their consumption baskets, ranging from

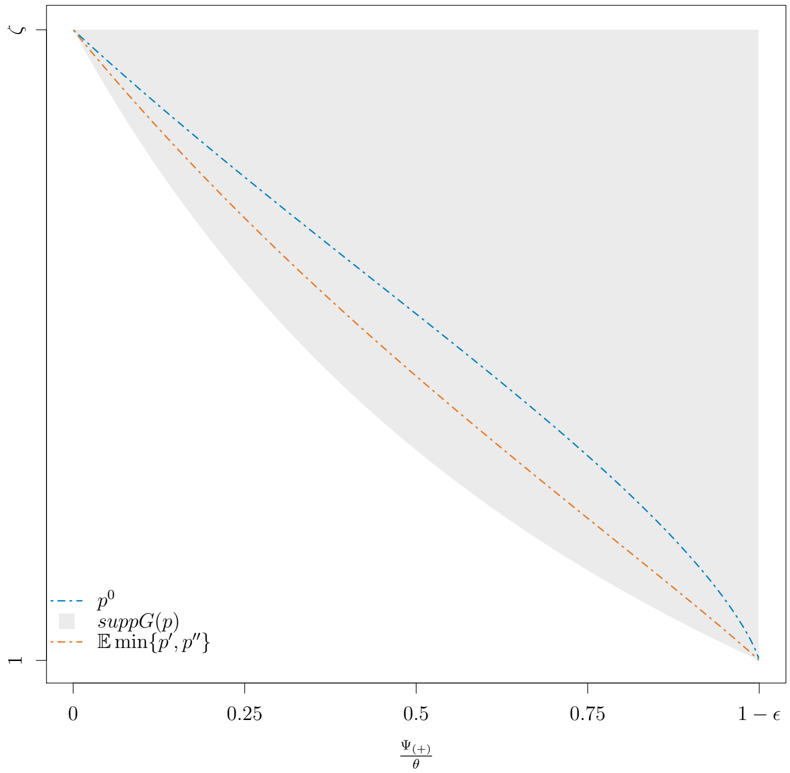

Figure 5: The figure illustrates the price support (the gray area), the expected price of a one-offer draw, p^0, and the expected minimum of two draws, $\mathbb{E}\min\{p', p''\}$, as functions of the average aggregate search intensity, $\frac{\Psi_{(+)}}{\theta} \in [0, 1)$. Given $\frac{\Psi_{(+)}}{\theta}$, a household can choose any price between $\mathbb{E}\min\{p', p''\}$ and p^0. This illustration highlights a significant search externality: the average aggregate price search intensity, constituted by the decisions of all households, influences the pricing problem faced by an individual household.

$\mathbb{E}\min\{p', p''\}$ to p^0.[31] Interestingly, due to search externalities, fully captive households in an economy populated by many deal-seeking customers may enjoy lower prices than a solitary bargain hunter in a predominantly captive-consumer environment.

[31] It is worth noting that $\mathbb{E}(p|s = 0) = p^0$, while $\lim_{s\to 1^-} \mathbb{E}(p|s) = p^0 - \overbrace{(\mathbb{E}\max\{p', p''\} - p^0)}^{\text{MPB}} = \mathbb{E}\min\{p', p''\}$. Considering $\max\{p', p''\} = \frac{p'+p''}{2} + |p' - p''|$ and $\min\{p', p''\} = \frac{p'+p''}{2} - |p' - p''|$, we find $\mathbb{E}\min\{p', p''\} + \mathbb{E}\max\{p', p''\} = \mathbb{E}p' + \mathbb{E}p'' = 2p^0$, leading to the latter equality.

D. Mapping the Model to Data

The model has been calibrated using both external and internal choices. The parameters set externally have been assigned standard values used in macroeconomic literature and are presented in Table 7. The parameters governing purchasing frictions $\{\phi, \omega, \zeta, \beta\}$ have been adjusted to reflect new empirical findings from Table 8.

Demographics. The model operates on an annual basis. Households enter the labor market at the age of 25, retire at 60, and die at 90. This implies $T_{work} = 35$ and $T = 65$.

Preferences. Consumption preferences are modeled using a CRRA specification, $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$. The parameter for elasticity of relative risk aversion $\frac{1}{\sigma}$ was set to 0.5. As a result of calibration, I established the parameters for shopping disutility as specified in Equation (13), with $\phi = .619$ and the relative weight at $\omega = 45.36$. The home production technology parameter is set to $\zeta = 162$. The discount factor β aims to generate an aggregate wealth-income ratio of 2.5, with the interest rate r set to .03, resulting in $\beta = .969$.

Income process. The income process includes both transitory $\{\varepsilon_t\}$ and permanent $\{\eta_t\}$ components. In accordance with the literature, the log variances of these shocks are set to $\sigma_\varepsilon^2 = .05$ and $\sigma_\eta^2 = .01$. The age-dependent deterministic component, κ_t, is approximated using a quadratic regression based on PSID data, as in Kaplan and Violante (2010). Upon retirement, households receive a social security income that reflects their last working-age period income, with a replacement rate *repl* (Guvenen and Smith, 2014; Berger et al., 2015).

Table 7: External choices

Parameter	Interpretation	Value
T_{work}	Age of retirement	35
T	Length of life	65
σ	Risk aversion	2.0
$repl$	Retirement replacement rate	.45
σ_ε^2	Variance of the transitory shock	.05
σ_η^2	Variance of the permanent shock	.01
r	Interest rate	.03

Table 8: Internal targets

Moment		Data	Model
Transaction prices:	top v. bottom decile	1.7	1.7
	rich work. v. poor work.	1.045	1.05
HH price index:	retirees v. poor work.	1	1.01
Saving-income ratio		2.5	2.5

V. Quantitative results

The calibrated version of the model provides deeper insights into the purchasing frictions involved, which are unattainable without a structural framework. In this section, I carry out two quantitative exercises. Firstly, I investigate the role of the price channel in influencing consumption expenditure changes due to not only transitory shocks but also other relevant household states. Secondly, through a simple counterfactual, I assess the extent of search frictions by determining the proportion of the population that benefits from being grouped with other consumers versus those who are disadvantaged by it.

A. Prices and Consumption Decisions: Model Perspective

In Subsection A.2 of the empirical section, I leveraged the quasi-experimental setup provided by the Economic Stimulus Act of 2008. This analysis identified the role of price adjustments in the overall response of household consumption expenditures in a specific local context. This quantitative exercise aims to extend that analysis further by employing the calibrated version of the model. Using artificial panel data from 10,000 households simulated by the model, I assess the significance of the price channel in influencing consumers' consumption decisions in the presence of the price dispersion.

To accomplish this, I leverage the simulated data to estimate the following equation:

$$\ln Z_{it} = \alpha_0 + \alpha_1 \varepsilon_{it} + \alpha_2 \nu_{it} + \alpha_3 \ln a_{it} + \xi_{it}, \tag{24}$$

where Z_{it} represents either total expenditures, $p_{it}c_{it}$, or prices, p_{it}, for households i at age t. In each specification, the dependent variable is regressed against the logs of income shocks and asset holdings, a_{it}.[32]

[32] Both income shocks, the transitory, ε_{it}, and the permanent, ν_{it}, are defined in Equation (15).

Utilizing the estimates and the decomposition outlined in Equation (6), I can dissect the total response of consumption expenditures into components attributable to income shocks and those due to price adjustments. Table 9 offers a quantitative summary of these findings. The results underscore the substantial influence of the price channel on household consumption decisions, with its importance being similar across various states. Intriguingly, the magnitude of its contribution is comparable to the lower bound of estimates from the empirical section, despite these estimates not being used in the model's calibration procedure.

Channel	Contribution
Assets (a_{it})	8.17%
Persistent income shocks (ν_{it})	8.13%
Transitory income shocks (ε_{it})	7.78%

Table 9: Using an artificial panel of 10,000 working-age households and the decomposition method from Equation (6), I estimate Equation (24) for prices and consumption expenditures. This approach allows me to report the contribution of the price channel from the responses of consumption expenditures to changes in assets, persistent income shocks, and transitory income shocks.

B. Price Search Externalities: Winners and Losers

Based on empirical evidence, the baseline model assumes that all consumers participate in one commonly shared market. As stated before, a key aspect from retailers' perspective is managing two types of uncertainty: not knowing if a customer will receive a competing price offer and being unaware of each visitor's specific probability of obtaining such offer.

In this exercise, I study a counterfactual outcome to see what happens if the latter layer of uncertainty were unveiled. Retailers still do not know whether a customer has received one or two price offers, but they are aware of the likelihood of two offers per customer. This results in price dispersion in each market, yet consumers encounter distinct conditions in their respective local goods market. The approach allows for evaluating the impact of price search externalities. By keeping households' shopping decisions unchanged, I explore the potential price distribution outcomes as indicated by Equation (23) from Theorem 1. In essence, this analysis involves setting up unique equilibrium price distributions for different types of households.[33] This method enables an examination of how price search externalities

[33] For each household type, x, the market is defined by $\Psi_{(-)} = c_t(x)(1 - s_t(x))$ and $\Psi_{(+)} = c_t(x)2s_t(x)$. The equilibrium price distribution (within its support) is then determined by $G(p) = \frac{1+s_t(x)}{2s_t(x)} - \frac{1-s_t(x)}{2s_t(x)} \cdot \frac{\zeta-1}{p-1}$.

influence market outcomes, highlighting the nuanced interplay between household behavior and market dynamics.

The analysis reveals that about 67% of households would benefit from being in separate markets, enjoying lower prices as a result. On the other hand, about 33% would prefer the shared market scenario, facing higher prices if markets were segmented. This insight underscores the significant role of market structure and consumer behavior in determining price levels across different segments of the economy.

VI. Concluding remarks

In this paper, I unveil new insights into shopping frictions amid household heterogeneity. I discover that price disparities across the income distribution are more pronounced than previously acknowledged and I establish a causal link between household income and the prices paid. Notably, the price channel contributes from 8 to 22% of consumption expenditure responses. Despite a wider variety in the consumption baskets of wealthier households, very few goods are specifically tailored to distinct income groups. This observation suggests significant price externalities, where the shopping behavior of one household influences another's. These insights are coherently explained by a novel incomplete-market model incorporating a random price-search mechanism.

Access to new, large datasets enabled researchers to study household consumption and shopping decisions with unprecedented granularity. The primary challenge in leveraging the NielsenIQ dataset for macroeconomic research lies in data sparsity. This paper significantly alleviates this issue, yet the need for further more systematic solution persists. From my perspective, future research could address the sparsity issue in two distinct ways: by employing specialized machine-learning methods designed for sparse datasets, or through more data-driven product aggregation. In another work (Pytka ⓡ Runge, 2024), we adopt the former strategy to examine consumption polarization. Alternatively, the latter approach could be explored using embedding techniques from the natural-language-processing literature.[34]

On the theoretical side, the paper demonstrates how to integrate purchasing frictions into a standard incomplete markets model in a manner that is both tractable and quantitatively significant. This model framework can be applied to a wide range of macroeconomic inquiries, particularly those where demand plays a more pronounced role in shaping economic

[34] Gentzkow et al. (2019a) provide an excellent survey of those methods from the economic perspective.

aggregates. One of key assumptions facilitating tractability of this framework is the transient nature of retailer-household matches. Future enhancements to this model should focus on establishing enduring relationships between retailers and consumers. Incorporating such persistence has the potential to introduce new truly microfounded real rigidities, offering a promising avenue for future research.

References

Acemoglu, Daron and Robert Shimer (1999) "Efficient unemployment insurance," *Journal of political Economy*, 107 (5), 893–928.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll (2021) "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 89 (1), 45–86, 10.1093/restud/rdab002.

Aguiar, Mark and Erik Hurst (2007) "Life-Cycle Prices and Production," *American Economic Review*, 97 (5), 1533–1559, http://ideas.repec.org/a/aea/aecrev/v97y2007i5p1533-1559.html.

_ , _ , and Loukas Karabarbounis (2013) "Time Use during the Great Recession," *American Economic Review*, 103 (5), 1664–96, 10.1257/aer.103.5.1664.

Aiyagari, S Rao (1994) "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109 (3), 659–84, http://ideas.repec.org/a/tpr/qjecon/v109y1994i3p659-84.html.

Argente, David and Munseob Lee (2020) "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 19 (2), 913–952, 10.1093/jeea/jvaa018.

Bai, Yan, José-Víctor Ríos-Rull, and Kjetil Storesletten (2024) "Demand Shocks as Technology Shocks," Working Paper 32169, National Bureau of Economic Research, 10.3386/w32169.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2015) "House Prices and Consumer Spending," Working Paper 21667, National Bureau of Economic Research, 10.3386/w21667.

Billingsley, Patrick (1995) *Probability and Measure*: Wiley, 3rd edition, http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0471007102.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008) "Consumption Inequality and Partial Insurance," *American Economic Review*, 98 (5), 1887–1921, 10.1257/aer.98.5.1887.

Broda, Christian, Ephraim Leibtag, and David E. Weinstein (2009) "The Role of Prices in Measuring the Poor's Living Standards," *Journal of Economic Perspectives*, 23 (2), 77–97, 10.1257/jep.23.2.77.

_ and Jonathan A. Parker (2014) "The Economic Stimulus Payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 68 (S), 20–36, 10.1016/j.jmoneco.2014.09.

Burdett, Kenneth and Kenneth L Judd (1983) "Equilibrium Price Dispersion," *Econometrica*, 51 (4), 955–69, http://ideas.repec.org/a/ecm/emetrp/v51y1983i4p955-69.html.

_ , Shouyong Shi, and Randall Wright (2001) "Pricing and matching with frictions," *Journal of Political Economy*, 109 (5), 1060–1085.

Butters, Gerard R. (1977) "Equilibrium Distributions of Sales and Advertising Prices," *The Review of Economic Studies*, 44 (3), 465–491, http://www.jstor.org/stable/2296902.

Chen, Fanglin, Xiao Liu, Davide Proserpio, and Isamar Troncoso (2022) "Product2Vec: Leveraging representation learning to model consumer product choice in large assortments," *NYU Stern School of Business*.

Commault, Jeanne (2022) "Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods," *American Economic Journal: Macroeconomics*, 14 (2), 96–122.

Diamond, Peter A. (1971) "A model of price adjustment," *Journal of Economic Theory*, 3 (2), 156–168, https://ideas.repec.org/a/eee/jetheo/v3y1971i2p156-168.html.

Dixit, Avinash K and Joseph E Stiglitz (1977) "Monopolistic competition and optimum product diversity," *The American economic review*, 67 (3), 297–308.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019a) "Text as data," *Journal of Economic Literature*, 57 (3), 535–574.

_ , Jesse M. Shapiro, and Matt Taddy (2019b) "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech," *Econometrica*, 87 (4), 1307–1340, https://doi.org/10.3982/ECTA16566.

Guerrieri, Veronica and Guido Lorenzoni (2017) "Credit crises, precautionary savings, and the liquidity trap," *The Quarterly Journal of Economics*, 132 (3), pp. 1427–1467, https://www.jstor.org/stable/26372706.

Guvenen, Fatih and Anthony A. Smith (2014) "Inferring Labor Income Risk and Partial Insurance From Economic Choices," *Econometrica*, 82 (6), 2085–2129, 10.3982/ECTA9446.

Handbury, Jessie (2021) "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," *Econometrica*, 89 (6), 2679–2715, https://doi.org/10.3982/ECTA11738.

Heathcote, Jonathan, Fabrizio Perri, Giovanni L. Violante, and Lichen Zhang (2023) "More unequal we stand? Inequality dynamics in the United States, 1967–2021," *Review of Economic Dynamics*, 50, 235–266, https://doi.org/10.1016/j.red.2023.07.014.

Hristakeva, Sylvia (2022) "Vertical Contracts with Endogenous Product Selection: An Empirical Analysis of Vendor Allowance Contracts," *Journal of Political Economy*, 130 (12), 3202–3252, 10.1086/720631.

Huggett, Mark (1993) "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 17 (5-6), 953–969, http://ideas.repec.org/a/eee/dyncon/v17y1993i5-6p953-969.html.

Huo, Zhen and José-Víctor Ríos-Rull (2015) "Tightening Financial Frictions on Households, Recessions, and Price Reallocations," *Review of Economic Dynamics*, 18 (1), 118–139, http://ideas.repec.org/a/red/issued/14-45.html.

Hwang, Minha, Bart J Bronnenberg, and Raphael Thomadsen (2010) "An empirical analysis of assortment similarities across US supermarkets," *Marketing Science*, 29 (5), 858–879.

Jaffe, Adam B. (1986) "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value," *The American Economic Review*, 76 (5), 984–1001, http://www.jstor.org/stable/1816464.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006) "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96 (5), 1589–1610, 10.1257/aer.96.5.1589.

Kang, ShinHyuck (2018) "Cyclical Dynamics of Shopping: Aggregate Implications of Goods and Labor Markets," *Available at SSRN 3283175*.

Kaplan, Greg and Guido Menzio (2015) "The Morphology of Price Dispersion," *International Economic Review*.

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

_ and Giovanni L. Violante (2010) "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2 (4), 53–87, 10.1257/mac.2.4.53.

Light, Bar (2020) "Uniqueness of equilibrium in a Bewley–Aiyagari model," *Economic Theory*, 69 (2), 435–450.

Mangin, Sephorah and Guido Menzio (2024) "Private and Social Learning in Frictional Product Markets,"Technical report.

Menzio, Guido (2023) "Search Theory of Imperfect Competition with Decreasing Returns to Scale," *Journal of Economic Theory* (forthcoming).

— and Shouyong Shi (2011) "Efficient Search on the Job and the Business Cycle," *Journal of Political Economy*, 119 (3), 468 – 510, http://ideas.repec.org/a/ucp/jpolec/doi10.1086-660864.html.

Michelacci, Claudio, Luigi Paciello, and Andrea Pozzi (2021) "The Extensive Margin of Aggregate Consumption Demand," *The Review of Economic Studies*, 10.1093/restud/rdab041.

Moen, Espen R (1997) "Competitive Search Equilibrium," *Journal of Political Economy*, 105 (2), 385–411, http://ideas.repec.org/a/ucp/jpolec/v105y1997i2p385-411.html.

Nord, Lukas (2023) "Shopping, Demand Composition, and Equilibrium Prices,"Technical report.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland (2013) "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103 (6), 2530–2553, https://ideas.repec.org/a/aea/aecrev/v103y2013i6p2530-53.html.

Pytka, Krzysztof (2017) *Macroeconomics of Bargain Hunting*, Ph.D. dissertation, European University Institute, https://hdl.handle.net/1814/49684.

— (2018) "Shopping Effort in Self-Insurance Economies,"Technical report.

— ⓡ Daniel Runge (2024) "Looking into the Consumption Black Box: Evidence from Scanner Data,"Technical report.

Rogerson, Richard (1988) "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, 21 (1), 3–16, http://ideas.repec.org/a/eee/moneco/v21y1988i1p3-16.html.

Ríos-Rull, José-Víctor (1995) *Frontiers of Business Cycle Research*, Chap. Models with Heterogeneous Agents.

— (1996) "Life-Cycle Economies and Aggregate Fluctuations," *The Review of Economic Studies*, 63 (3), 465–489, http://www.jstor.org/stable/2297891.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod (2010) "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," *Tax Policy and the Economy*, 24 (1), 69–110, 10.1086/649829.

Stigler, George J. (1961) "The Economics of Information," *Journal of Political Economy*, 69, 213, http://ideas.repec.org/a/ucp/jpolec/v69y1961p213.html.

Uhlig, Harald (1996) "A Law of Large Numbers for Large Economies," *Economic Theory*, 8 (1), 41–50.

Online Appendix

A. Representativeness of the dataset

In this section, I examine the sensitivity of the dataset to the restrictions applied in the analysis. As demonstrated, focusing on households whose current income information is available does not substantially affect the implied dynamics (Figure 6) of aggregate consumption nor the household composition of the dataset (Table 10). Conversely, excluding products that are not purchased frequently enough results in dynamics markedly different from both the unrestricted NielsenIQ dataset and non-durable consumption in the Consumption Expenditure Survey (CEX).[1] For this reason, I decide to use the dataset with transactions for all products and for households for whom there is information on the current income.

Table 10: Fraction of households by certain characteristics. Column *NielsenIQ* presents statistics for the unrestricted dataset, while *NielsenIQ Current* provides a breakdown for households with information on their current income.

Variable	NielsenIQ	NielsenIQ Current
Married	0.45	0.43
No Children	0.66	0.71
Female Head Employed	0.58	0.57
Male Head Employed	0.71	0.68
Female High School	0.40	0.42
Male High School	0.39	0.39
White	0.77	0.77

[1] As Heathcote et al. (2023) demonstrate, despite the considerable gap in expenditure levels between the CEX and the NIPA, the growth in CEX expenditures over the most recent period (2004-2018) closely tracks the growth observed in NIPA, more accurately than the PSID. In this context, the business cyclicality of the CEX is considered as the benchmark for dynamic representativeness in this section.

Figure 6: Dynamics of implied aggregate consumption for different dataset restrictions. *NielsenIQ* represents the dynamics for the unrestricted dataset. *NielsenIQ 10* depicts the consumption dynamics after excluding products purchased 10 times or less in a given month, focusing on more frequently purchased items. *NielsenIQ Current* refers to the consumption for households for whom current income information was available (utilizing data from future waves). For reference, the implied consumption from the CEX is also included to provide a benchmark comparison. The reference year is 2004. All measures in the NielsenIQ dataset are calculated using projection factors.

B. Hand-to-Mouth households pay lower prices

An additional intriguing aspect of the heterogeneity in prices paid is its variation with different levels of financial liquidity. To explore this, I utilize a tax rebates survey conducted by NielsenIQ for Broda and Parker (2014) in 2008. This survey posed the question to panelists:

> *In the event of an unexpected decline in income or increase in expenses, do you have at least two months of income available in cash, bank accounts, or easily accessible funds?*

The responses serve as a proxy for identifying hand-to-mouth households, characterized by insufficient liquidity.

Table 11 shows regression estimates for households responding to the liquidity question, including two new variables: the liquidity state and the interaction between high earnings and the liquidity state. It appears that constrained households, not earning above the median level, pay between 0 and 2.8% lower prices. The effect for high-earning hand-to-mouth households is smaller, or in some cases, entirely neutralized. These findings indicate that the price indices of consumption bundles purchased by households are influenced not only by current labor income but also by the household's balance sheet.

Table 11: Household price indices and financial liquidity

	$\ln \bar{P}_{j,m}$			
	(1)	(2)	(3)	(4)
HH:HtM	-0.003	-0.007^{***}	-0.028^{***}	-0.024^{***}
	(0.002)	(0.002)	(0.004)	(0.003)
HH:HtM & HH Earnings > median(HH Earnings)	0.004	0.004	0.015^{**}	0.013^{**}
	(0.004)	(0.003)	(0.007)	(0.005)
HH Earnings > median(HH Earnings)	0.021^{***}	0.015^{***}	0.071^{***}	0.054^{***}
	(0.002)	(0.002)	(0.004)	(0.003)
Non-employed in working age (Female)	-0.009^{***}	-0.008^{***}	-0.014^{**}	-0.010^{**}
	(0.003)	(0.003)	(0.006)	(0.004)
Non-employed in working age (Female)	-0.008^{***}	-0.004^{**}	-0.007^{*}	-0.004
	(0.002)	(0.002)	(0.004)	(0.003)
Retired (Male)	-0.009	-0.004	-0.007	-0.003
	(0.006)	(0.005)	(0.011)	(0.008)
Retired (Female)	0.009^{*}	0.007	0.018^{*}	0.013^{*}
	(0.006)	(0.004)	(0.010)	(0.008)
HH composition dummies	*Yes*	*Yes*	*Yes*	*Yes*
Age dummies (both heads)	*Yes*	*Yes*	*Yes*	*Yes*
Month dummies	*Yes*	*Yes*	*Yes*	*Yes*
Year dummies	*Yes*	*Yes*	*Yes*	*Yes*
Scantrack market dummies	*Yes*	*Yes*	*Yes*	*Yes*
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	284,112	284,112	284,112	284,112
Number of panelists	24,141	24,141	24,141	24,141
R^2	0.040	0.020	0.088	0.056

Notes: Robust standard errors clustered at the household level are included in parentheses.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

C. A Minor Role of Store Amenities

To examine the impact of store amenities, I constructed a revenue-weighted average of the prices of different goods sold at store s in month m, following the methodology proposed by Kaplan and Menzio (2015, section 3.2), denoted as $\mu_{s,m}$. Utilizing these indices, I then calculated the quantity-weighted expensiveness level for each household j:

$$\bar{\mu}_{j,m} = \frac{\sum_s \mu_{s,m} \cdot \text{exp}_{s,j,m}}{\sum_s \text{exp}_{s,j,m}}, \tag{25}$$

where $\text{exp}_{s,j,m}$ represents the expenditures made by household j in store s during month m. The analysis, as detailed in Table 12, reveals that the variation in $\ln \bar{\mu}_{j,m}$ is notably low. Specifically, the values of the store-specific component estimates range from 17 to 30% of the variation observed in the price regression from Table 2. This finding underscores the relatively minor role that store-specific pricing plays in the overall heterogeneity of household price indices.

Table 12: Store expensiveness across different income and employment state

	$\ln \bar{\mu}_{j,m}$			
	(1)	(2)	(3)	(4)
HH Earnings > median(HH Earnings)	0.006***	0.004***	0.013***	0.009***
	(0.0004)	(0.0002)	(0.001)	(0.0004)
Non-employed in working age (Male)	−0.001	−0.0002	−0.001	−0.001
	(0.001)	(0.001)	(0.002)	(0.001)
Non-employed in working age (Female)	−0.0003	−0.0004	−0.003**	−0.003**
	(0.001)	(0.001)	(0.001)	(0.001)
Retired (Male)	−0.001	−0.001***	−0.002**	−0.003***
	(0.001)	(0.0004)	(0.001)	(0.001)
Retired (Female))	−0.002***	−0.001***	−0.001	−0.001**
	(0.0004)	(0.0003)	(0.001)	(0.0005)
HH composition dummies	Yes	Yes	Yes	Yes
Age dummies (both heads)	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Scantrack market dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	4,751,339	4,751,201	4,751,395	4,691,551
Number of panelists	91,150	91,156	91,142	91,150

D. Decomposition of responses to tax returns

Tables 13 and 14 present the estimation results. Table 14, focusing on reactions to $Q_{j,m}$, displays results for four considered definitions of goods. Table 13, which reports the pass-through to overall expenditures, features only one specification. This is because $\bar{P}_{j,m}Q_{j,m}$, by construction, does not vary across definitions. The estimated overall response of total expenditures to the receipt of the ESP amounts to approximately 5.5-7.5% of pre-treatment consumption, with evidence suggesting an anticipatory response just before receipt. These findings align with a similar analysis conducted by Michelacci, Paciello, and Pozzi (2021).[2] The reactions in $Q_{j,m}$ are notably smaller than those in $\bar{P}_{j,m}Q_{j,m}$, attributable to the positive price adjustments observed in Table 3 and Figure 7.

Table 13: Expenditure response to the ESP

Response to the ESP	$\ln\left(\bar{P}_{j,m}Q_{j,m}\right)$
Quarter before, β_{-1}	0.024***
	(0.008)
Quarter of receipt, β_0	0.054***
	(0.012)
One quarter after, β_1	0.073***
	(0.016)
Two quarters after, β_2	0.076***
	(0.020)
Month dummies	Yes
Number of observations	345,768
Number of panelists	29,289
R^2	0.659

Notes: ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

[2] Interestingly, the authors utilize the same data sources to examine the total expenditure response to the ESP. Despite minor technical differences, the principal distinction lies in their focus on changes in the products entering households' consumption baskets. Conversely, my analysis zeroes in on the prices paid by households for identical or very similar products, an aspect not covered by Michelacci, Paciello, and Pozzi (2021).

Table 14: Consumption response to the ESP

Response to the ESP	$\ln Q_{j,m}$			
	(1)	(2)	(3)	(4)
Quarter before, β_{-1}	0.022***	0.023***	0.022***	0.021***
	(0.008)	(0.008)	(0.008)	(0.008)
Quarter of receipt, β_0	0.048***	0.050***	0.045***	0.046***
	(0.012)	(0.012)	(0.012)	(0.012)
One quarter after, β_1	0.066***	0.068***	0.064***	0.063***
	(0.016)	(0.016)	(0.016)	(0.016)
Two quarters after, β_2	0.068***	0.070***	0.065***	0.063***
	(0.020)	(0.020)	(0.019)	(0.019)
Month dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	345,768	345,768	345,768	345,768
Number of panelists	29,289	29,289	29,289	29,289
R^2	0.661	0.660	0.656	0.656

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

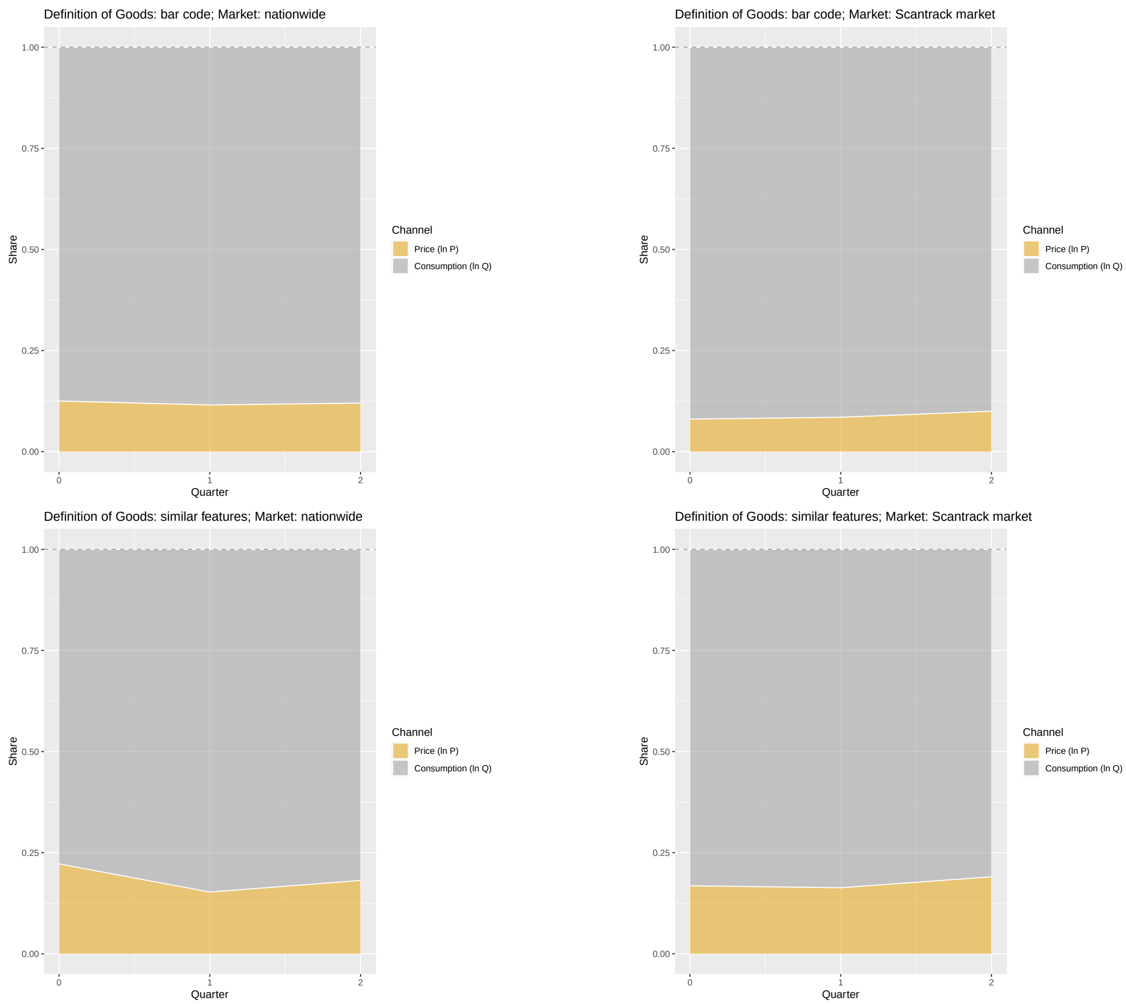

Figure 7: Decomposition (6) of the expenditure responses to the ESP

To decompose the differences in expenditures, $\mathbb{E}[\bar{P}_{j,\tau+s}Q_{j,\tau+s} - \bar{P}_{j,\tau-1}Q_{j,\tau-1}]$, rather than the expenditure growth, $\mathbb{E}\ln\left(\frac{\bar{P}_{j,\tau+s}Q_{j,\tau+s}}{\bar{P}_{j,\tau-1}Q_{j,\tau-1}}\right)$, I proceed as follows:

$$\mathbb{E}[\bar{P}_{j,\tau+s}Q_{j,\tau+s} - \bar{P}_{j,\tau-1}Q_{j,\tau-1}] = \mathbb{E}[\bar{P}_{j,\tau+s}Q_{j,\tau+s}] - \mathbb{E}[\bar{P}_{j,\tau-1}Q_{j,\tau-1}].$$

Using the definition of the covariance we have:

$$\begin{aligned}
\mathbb{E}[\bar{P}_{j,\tau+s}Q_{j,\tau+s}] &= \mathbb{E}[\bar{P}_{j,\tau+s}]\mathbb{E}[Q_{j,\tau+s}] + \mathrm{Cov}(\bar{P}_{j,\tau+s}, Q_{j,\tau+s}),\\
\mathbb{E}[\bar{P}_{j,\tau-1}Q_{j,\tau-1}] &= \mathbb{E}[\bar{P}_{j,\tau-1}]\mathbb{E}[Q_{j,\tau-1}] + \mathrm{Cov}(\bar{P}_{j,\tau-1}, Q_{j,\tau-1}).
\end{aligned}$$

Next, observe that:

$$\begin{aligned}
&\mathbb{E}[\bar{P}_{j,\tau+s}]\mathbb{E}[Q_{j,\tau+s}] - \mathbb{E}[\bar{P}_{j,\tau-1}]\mathbb{E}[Q_{j,\tau-1}] =\\
&\qquad = \mathbb{E}[Q_{j,\tau+s}] \cdot \Big(\mathbb{E}[\bar{P}_{j,\tau+s}] - \mathbb{E}[\bar{P}_{j,\tau-1}]\Big) + \mathbb{E}[\bar{P}_{j,\tau-1}] \cdot \Big(\mathbb{E}[Q_{j,\tau+s}] - \mathbb{E}[Q_{j,\tau-1}]\Big) =\\
&\qquad = \mathbb{E}[Q_{j,\tau+s}] \cdot (1 - L^{k+1})\mathbb{E}[\bar{P}_{j,\tau+s}] + \mathbb{E}[\bar{P}_{j,\tau-1}] \cdot (1 - L^{k+1})\mathbb{E}[Q_{j,\tau+s}],
\end{aligned}$$

where L^k is the lag operator of order k. Finally we have:

$$\begin{aligned}
\mathbb{E}[\bar{P}_{j,\tau+s}Q_{j,\tau+s} - \bar{P}_{j,\tau-1}Q_{j,\tau-1}] = &\overbrace{\mathbb{E}[Q_{j,\tau+s}] \cdot (1 - L^{s+1})\mathbb{E}\left[\bar{P}_{j,\tau+s}\right]}^{\text{Price channel}} + \overbrace{\mathbb{E}\bar{P}_{j,\tau-s} \cdot (1 - L^{s+1})\mathbb{E}[Q_{j,\tau+s}]}^{\text{Consumption channel}} +\\
&+ \underbrace{(1 - L^{s+1})\mathrm{Cov}\left(\bar{P}_{j,\tau+s}, Q_{j,\tau+s}\right)}_{\text{Interactions}}. \qquad (26)
\end{aligned}$$

All quantities necessary for the decomposition are calculated and presented in Tables 15, 16, and 17. The decomposition itself is presented in Figure 8.

Table 15: Consumption response to the ESP

Response to the ESP	$\bar{P}_{j,m}$			
	(1)	(2)	(3)	(4)
Q_1	0.001	0.001	0.002	0.002
	(0.002)	(0.001)	(0.002)	(0.002)
Q0	0.005*	0.003	0.008**	0.006**
	(0.003)	(0.002)	(0.004)	(0.003)
Q1	0.009**	0.005*	0.010*	0.010**
	(0.004)	(0.003)	(0.005)	(0.004)
Q2	0.009**	0.007*	0.012*	0.011**
	(0.004)	(0.004)	(0.006)	(0.005)
Month dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	345,768	345,768	345,768	345,768
Number of panelists	29,289	29,289	29,289	29,289
R^2	0.441	0.434	0.505	0.492

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 16: Expenditure response to the ESP

Response to the ESP	$Q_{j,m}$			
	(1)	(2)	(3)	(4)
Q_1	6.577***	6.642***	7.007***	6.760***
	(2.309)	(2.331)	(2.288)	(2.308)
Q0	20.037***	20.695***	19.640***	19.787***
	(3.489)	(3.543)	(3.429)	(3.477)
Q1	25.085***	25.786***	25.341***	24.500***
	(4.482)	(4.542)	(4.395)	(4.458)
Q2	24.787***	25.788***	24.762***	24.272***
	(5.426)	(5.526)	(5.306)	(5.420)
Month dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	345,768	345,768	345,768	345,768
Number of panelists	29,289	29,289	29,289	29,289
R^2	0.727	0.723	0.730	0.725

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 17: Price response to the ESP

	$\bar{P}_{j,m}Q_{j,m}$			
	(1)	(2)	(3)	(4)
Q_1	7.845***	7.845***	7.845***	7.845***
	(2.374)	(2.374)	(2.374)	(2.374)
Q0	23.625***	23.625***	23.625***	23.625***
	(3.681)	(3.681)	(3.681)	(3.681)
Q1	28.908***	28.908***	28.908***	28.908***
	(4.747)	(4.747)	(4.747)	(4.747)
Q2	28.968***	28.968***	28.968***	28.968***
	(5.765)	(5.765)	(5.765)	(5.765)
Month dummies	Yes	Yes	Yes	Yes
Product aggregation	Bar code	Bar code	Features	Features
Area aggregation	Nationwide	Scantrack	Nationwide	Scantrack
Number of observations	345,768	345,768	345,768	345,768
Number of panelists	29,289	29,289	29,289	29,289
R^2	0.715	0.715	0.715	0.715

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

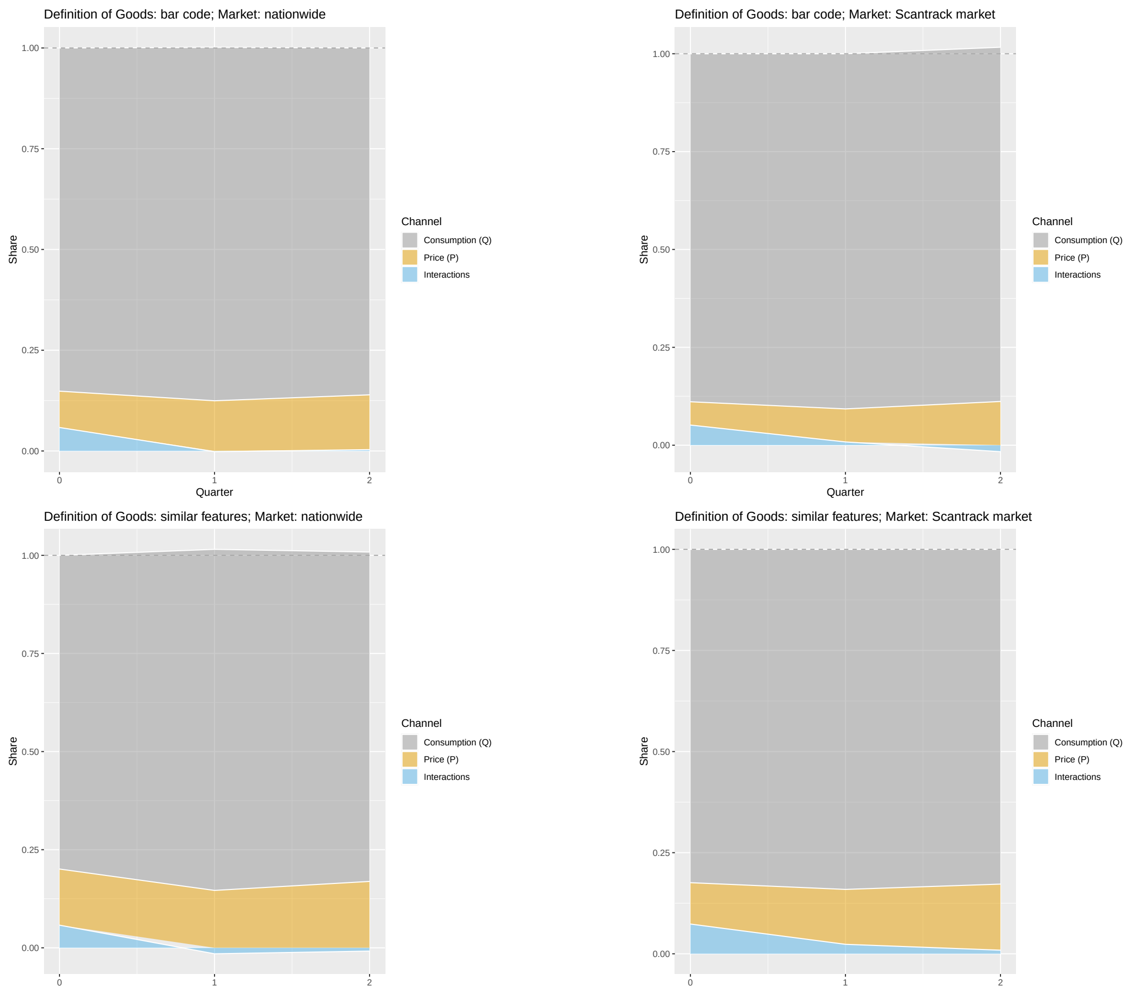

Figure 8: Decomposition (26) of the expenditure responses to the ESP

E. Finite-Sample Bias in Histogram Overlap

In a recent study, Nord (2023) analyzes the same dataset with a focus on multiproduct markets, where goods are customized for distinct consumer segments. The calibration of the model is driven by the reported substantial differences in consumption baskets between low- and high-expenditure households. To demonstrate this, they propose the following procedure: first, households from the NielsenIQ panel are divided into quintiles based on their consumption expenditure; next, the expenditure share that quintile g spends on good j in a given year is calculated:

$$\omega_j^g = \frac{\sum_{i \in g} e_j^i}{\sum_{j \in J} \sum_{i \in g} e_j^i}. \tag{27}$$

Given the distribution of spending across product dimensions for each household group, they propose a similarity measure called 'histogram overlap' to compare the top and bottom quintile groups:

$$\Omega^{g,h} = \sum_{j \in J} \min \left\{ \omega_j^g, \omega_j^h \right\}. \tag{28}$$

Any deviation from $\Omega^{g,h} = 1$ can be interpreted as evidence of existing products being purchased disproportionately more often by one household group. The author demonstrates that, although the similarity between the highest and lowest expenditure quintiles is relatively high (with an overlap of 86%) at a more general product definition (module level), it drops to only 63% at a more granular level (barcode level). This is interpreted as evidence of significant consumption polarization, and the result is used to calibrate the preferences in the quantitative model.

The main empirical conclusion of Nord (2023) is quite different from my findings from Subsection B from Section III. As I show, the discrepancy is likely to be caused by a strong finite-sample bias of the histogram overlap from Equation (28) in the used dataset, which comes from the inherent high-dimensionality of consumption.[3] Intuitively, there are many products (for example, $\approx 800,000$ unique products were purchased in 2014), but only slightly

[3] As outlined in Section F of the Appendix, my analysis of extensive and intensive margins is substantially less influenced by this bias. Additionally, similarity measures such as weighted correlation and cosine similarity are notably high. If at all, finite-sample bias would imply that the actual similarity measures could be even higher.

above 10,000 households per quintile. Many products will therefore be purchased only by a small number of households, potentially just one. A naïve estimator, such as the histogram overlap in Equation (28), without adjustments for data sparsity, may interpret products bought extremely rarely as evidence of polarization—even if both types of households selected the product with equal probability.

This bias has been recently discussed in more detail by Gentzkow, Shapiro, and Taddy (2019b). Although they study a very different problem, the polarization of US politics using congressional speech data, they encounter the same situation where the set of choices (two-word phrases that a speaker could use in their case, products that a consumer could buy in mine) is very large relative to the amount of choices actually observed. They develop a model of a speech-generating process, then analyze the bias more formally, and introduce estimators that can overcome finite-sample bias to recover a more accurate estimate of polarization.

Following a procedure similar to that proposed by Gentzkow et al. (2019b), to demonstrate that the bias affects the results, we need to document two aspects: first, that the overlap among infrequently purchased products is lower, and second, that a significant share of expenditure is allocated to these products. The former aspect shows that some observations are influenced by this bias, while the latter indicates the bias's quantitative significance.[4] Figure 9a shows that both elements are present in the data. For products which are purchased by less than 1% of households in at least one group elements of the histogram overlap, $\min\{\omega_j^g, \omega_j^h\}$, are extremely low, but they rapidly increase and are almost constant above 80% for all products which are purchased by a share of 1% or more of households in both groups. The expenditure share is also heavily concentrated on infrequently purchased goods, with goods being purchased by less than 1% of households in a group making up 80% of total expenditure.

While these results indicate that the bias is quite severe, we cannot quantify its extent. Since there might be systematic differences in how both groups purchase these infrequently purchased items, we cannot identify which share of the reduction in overlap to assign to the bias and which to actual differences in consumption patterns. To overcome this problem Gentzkow et al. (2019b) suggest a permutation test. In the context of the considered problem, each household is randomly assigned to a 'placebo quantile'. Since this assignment is random, we know that the consumption probabilities and expected prices paid are identical between

[4]If almost all spending of the typical household would fall on frequently purchased items, the bias might still be present but would only have a minimal effect on the estimated overlap.

Figure 9: Analysis of Overlap by Purchasing Frequency and Placebo Test Results

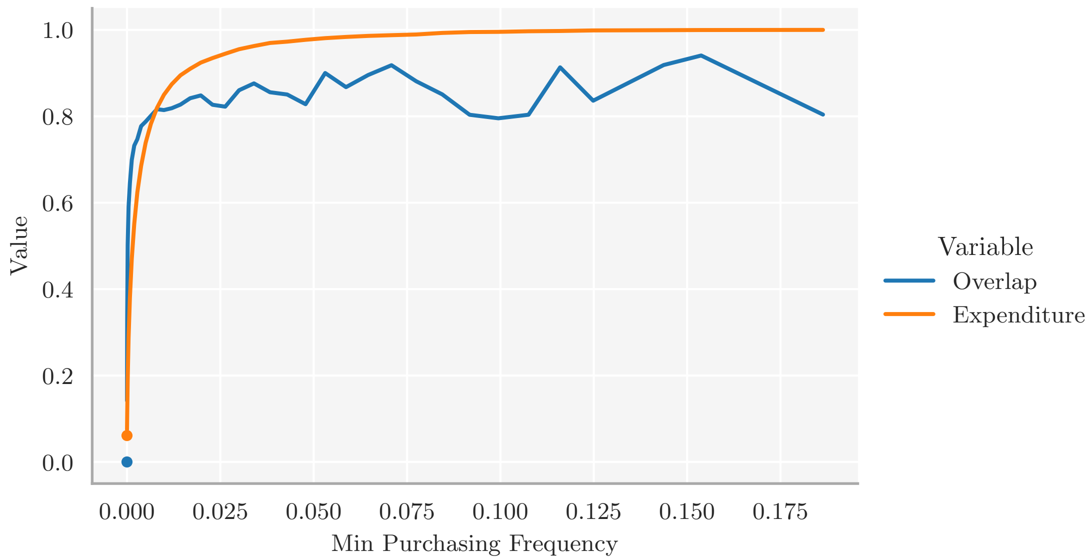

(a) Overlap by Purchasing Frequency: The figure plots the cumulative share of total expenditure (orange line) and the average overlap (blue line), conditioned on a specific minimum purchasing frequency between the first and last expenditure quintile in 2014, excluding magnet products. Dots represent products not purchased by either of the two groups, accounting for 6% of total expenditure. Products are ordered by the lowest frequency with which they are purchased by households in both groups, meaning that a minimum purchasing frequency of 0.1 indicates that a product is purchased by at least 10% of households in both expenditure quintiles.

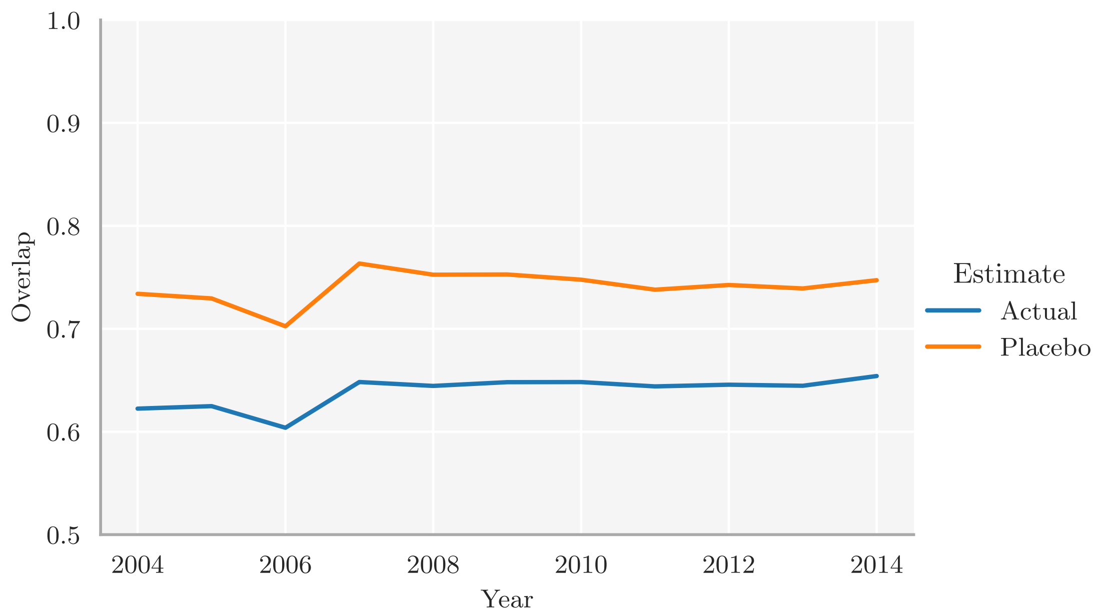

(b) Placebo test: The figure plots the basket overlap measure between the lowest and highest quintile of consumption expenditure (blue line). The placebo estimates (orange line) are created by randomly selecting households into 5 groups of equal size and then performing the overlap analysis using this placebo quantile.

the two groups and therefore in the absence of the finite-sample bias, the estimated overlap should be 1. Any deviation from that number can therefore be attributed to the finite-sample bias. Figure 9b shows the results of this test. While the actual estimate is still significantly lower than the one based on placebo quintiles, the figure still indicates that almost three-quarter of the effect documented by Nord (2023) is due to the finite-sample bias.

In a companion paper (Pytka ⓡ Runge, 2024), we study polarization more formally than in Subsection B from Section III. We estimate separate consumption models for households in the top and bottom quintiles using penalized multinomial models, which are designed to address data sparsity. These estimated consumption generating processes are then compared using polarization measures proposed by Gentzkow et al. (2019b). Our results closely mirror those presented in this paper, indicating an absence of polarization.

F. Finite-Sample Bias in the Decomposition into Extensive and Intensive Margins

While the decomposition into margins from Equation (7) could also be subject to finite-sample bias, similar to the histogram overlap discussed earlier, the direction of this bias remains unclear. To examine this potential bias, I conduct a similar analysis as presented in Section E of the appendix. Figure 10a illustrates the average contribution of the extensive margin, conditional on the minimum purchasing frequency observed between the two groups. Although the contribution shows extreme values for products purchased only by one group or purchased extremely infrequently, the two clearly visible outliers together only account for 4% of total expenditure; a further 21% of total consumption the values vary between 0.4 and 3.1. The vast majority of expenditure falls on products where the average contribution is quite close to the full panel estimate of 1.13, which indicates that the bias has a much lower impact than in the case of the histogram overlap estimator studied above.

When looking into these early outliers it is important to bear in mind that the plotted contribution is the ratio of two conditional averages, namely the average extensive margin divided by the average difference in consumption conditional on a product being purchased at a specific rate. These local estimators have two important consequences. First, if the finite-sample bias were to introduce random noise in this region that dominates any systematic effect, both of these conditional estimators would vary closely around zero, which will increase the variance of their ratio.[5] Similarly, all negative estimates of the contribution come from the fact that the products at these frequencies are purchased more by poorer households, which yields a negative quantity difference. Since the full-sample estimates do not rely on these local estimators, but instead on the ratio of global estimates, these issues are averaged out and no longer relevant.

To further support this result, I follow a procedure by Gentzkow et al. (2019b) and analogous to the one conducted for the histogram overlap in Section E of the appendix. Figure 10b shows the decomposition per year for both the actual data and a randomly assigned grouping. The results are quite striking. For the actual data we see that the extensive margin is above the quantity difference in all years, with a relatively constant difference between

[5] Some evidence that points towards this is that at the spike, the average quantity difference is only 1% of the average absolute difference in quantity. The values for later points are typically much larger indicating that a systematic effect is much stronger there.

Figure 10: Decomposition by Purchasing Frequency and Extensive-Margin Permutation Test

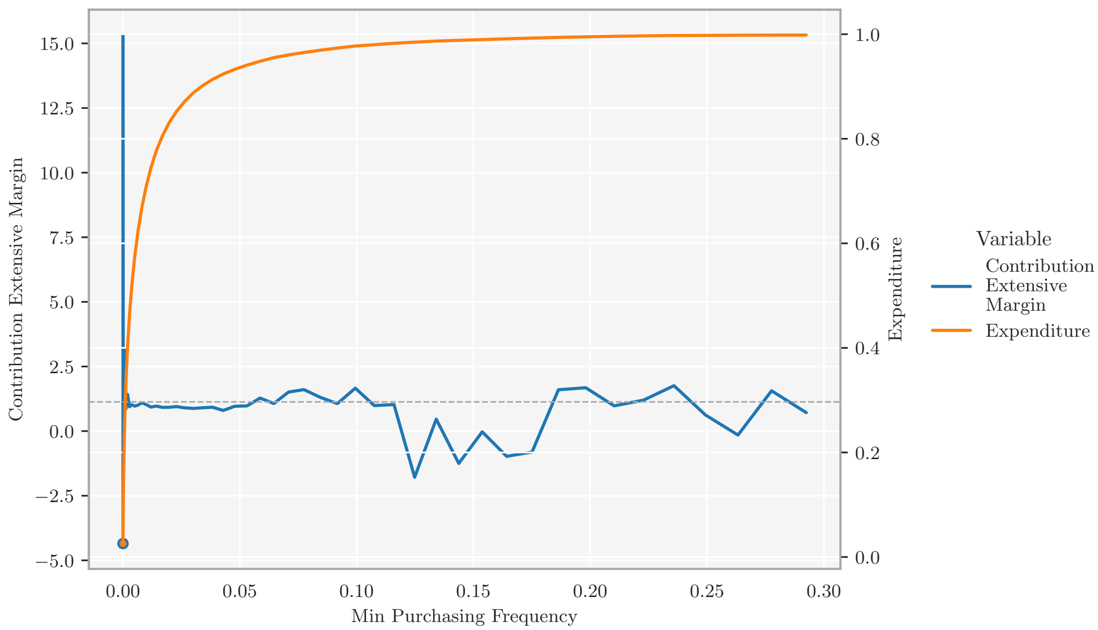

(a) Decomposition by Purchasing Frequency: The figure plots the cumulative share of total expenditure (orange line) and the average value of the extensive margin (blue line), conditional on a specific minimum purchasing frequency in both groups. The dashed line indicates the average contribution of 1.13.

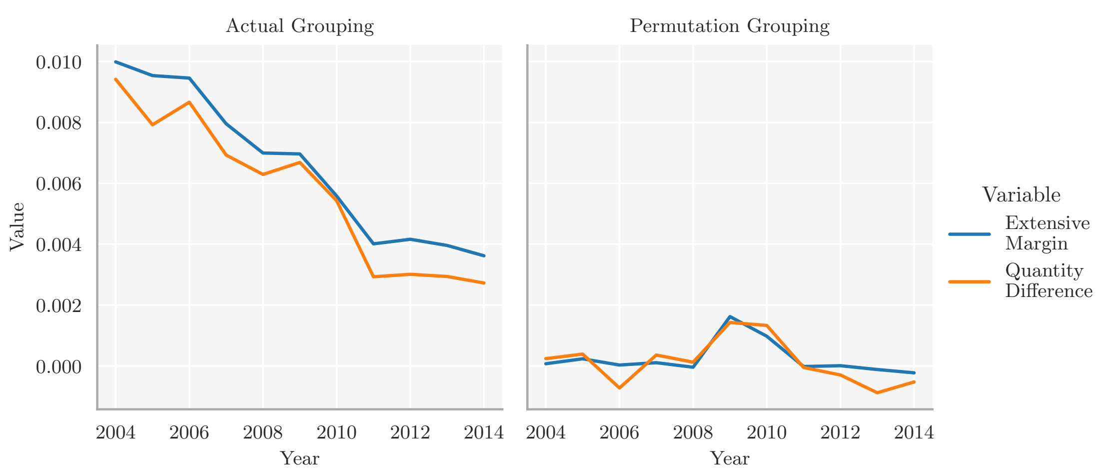

(b) Extensive-Margin Permutation Tests: The left panel shows the average extensive margin and the average consumption difference between rich and poor households, while the right panel shows the average extensive margin and average consumption difference between two placebo groups.

the two. On the other hand, the randomly assigned grouping behaves very differently, with both series generally varying close to zero with no systematic difference between them. This shows that the finite-sample bias creates primarily random noise, but does not affect either the extensive margin or the difference in consumption in any systematic way and therefore the relatively stable contribution shown in the left panel is due to systematic differences in purchasing behavior between rich and poor households.

G. Well-being of Shopping: High Earners Do Not Enjoy Shopping More

In this section I document whether there is a significant difference in perception of shopping across different individuals. To this end, I use the ATUS Well-Being Module, where respondents answer what they feel during reported activities. In particular, I am interested in the answers related to shopping and whether different employment groups experience this activity in a different way.

Data. The ATUS Well-Being (WB) Module is a complementary survey conducted by the U.S. Bureau of Labor Statistics in three waves 2010, 2012, and 2013. In this survey respondents are asked to evaluate their subjective well-being during reported activities. Those questions relate to experienced happiness, sadness, tiredness, stress, pain of activities. In all three waves there are over 75,000 respondents.

In the WB module households report happiness, sadness, tiredness, stress, pain of activities on a scale from 0 (not at all) to 6 (completely). I regressed the reported answers on dummies indicating: respondents with the total annual income above the median, non-employment status, retirement status, the reported activity is shopping (defined as in the previous subsection), and the interaction of the shopping activity and employment status of the individual. In addition to this I controlled for age categories, 'shopping needs', and time dummies: year, day, and daytime.

Table 18 presents the results of the estimation. As can be seen, there is no significant difference (at the significance level 5%) in well-being experienced while shopping across different groups. Those results allow to rule out the possibility that for some groups shopping is non-market work, while for others it is more like a leisure activity.

Table 18: Well-being, shopping, and employment status

	WUTIRED (1)	WUHAPPY (2)	WUPAIN (3)	WUSTRESS (4)	WUSAD (5)
Activity:Shopping & Earnings>median(Earnings)	−0.084	0.044	−0.176*	−0.154	−0.082
	(0.106)	(0.093)	(0.092)	(0.100)	(0.077)
Activity:Shopping & Nonemployed (in working age)	0.031	−0.108	−0.242***	0.189**	−0.110
	(0.092)	(0.080)	(0.080)	(0.087)	(0.067)
Activity:Shopping & Retired	0.055	−0.014	−0.230**	0.076	−0.046
	(0.112)	(0.098)	(0.097)	(0.106)	(0.082)
Activity:Shopping	−0.256***	0.052	0.010	−0.033	−0.028
	(0.069)	(0.060)	(0.060)	(0.065)	(0.050)
Earnings>median(Earnings)	0.017	−0.108***	−0.354***	0.238***	−0.171***
	(0.022)	(0.019)	(0.019)	(0.021)	(0.016)
Nonemployed (in working age)	0.014	−0.104***	0.558***	0.090***	0.192***
	(0.020)	(0.018)	(0.018)	(0.019)	(0.015)
Retired	−0.571***	0.188***	0.513***	−0.684***	0.047**
	(0.031)	(0.027)	(0.027)	(0.029)	(0.022)
Age categories	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Shopping needs	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year and day dummies	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Daytime dummy and duration control	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
N	76,506	76,506	76,506	76,506	76,506
R^2	0.052	0.020	0.068	0.054	0.025

Notes:

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

H. Proofs

A. Proof of Proposition 1

Proof. To derive (12) I use the fact that the expected value of any non-negative random variable x distributed according to a cdf $H(x)$ can be computed integrating over its survival function (Billingsley, 1995, p. 79), namely:

$$\mathbb{E}(x) = \int_0^\infty (1 - H(x))dx. \tag{29}$$

The price of the consumption bundle is then a result of applying this property to equation (10):

$$\mathbb{E}(p|s_t) = \int_0^\infty 1 - G(x) - s_t \left(G(x) - [G(x)]^2\right) dx,$$

where $\int_0^\infty 1 - G(x)dx$ is the expected value for the captive offer and, using an analogous reasoning from Lemma 1, is also the price of consumption for the fully captive household that decides not to make any search for prices.

The residual part is equal to:

$$\int_0^\infty \left(G(x) - [G(x)]^2\right) dx =: MPB, \tag{30}$$

and which is clearly always positive as $\forall_x G(x) \geq [G(x)]^2$. For better interpretation it is convenient to reformulate equation (30):

$$\int_0^\infty \left(G(x) - [G(x)]^2\right) dx = \underbrace{\int_0^\infty 1 - [G(x)]^2 dx}_{\mathbb{E}\max\{p',p''\}} - \underbrace{\int_0^\infty 1 - G(x)dx}_{p^0}.$$

The expected maximum of two independent draws, $\max\{p', p''\}$ is distributed according to $[G(x)]^2$. It can be easily shown by the fact that $Pr(\max\{p', p''\} \leq x) = Pr(p' \leq x, p'' \leq x)$. Assuming independence of p' and p'' we get $Pr(p' \leq x) \cdot Pr(p'' \leq x) = [G(x)]^2$. Therefore, $\mathbb{E}\max\{p', p''\} = \int_0^\infty 1 - [G(x)]^2 dx$.

□

B. Proof of Lemma 2

The two first properties are an immediate result of Lemma 1 from (Burdett and Judd, 1983). Suppose that $G(p)$ has a discontinuity at some $p' \in \text{supp } G(p)$. The retailer posting an infinitesimally smaller price $p' - \epsilon$ would increase its profit as the probability of making a sale would change by a discrete amount. Furthermore, supp $G(p)$ is a connected set. Suppose there is a gap of zero probability between p' and p''. The seller's gain would be strictly higher at p'' as $p'' > p'$, and $G(p') = G(p'')$. This cannot occur in an equilibrium.

Next, suppose that (iii) is not true. Then $\max \text{supp } G(p) =: \overline{p} \leq \zeta$.[6] Moreover, we know that $G(\overline{p}) = G(\zeta) = 1$. If we substitute values of the c.d.f. for both prices into (17) all firms will have incentives to set higher price for higher demand, which leads us to contradiction. As a result, $\max \text{supp } G(p) = \zeta$. Fact (iv) is an equilibrium condition. If there would be such a price p that would yield higher profit, each individual retailer would have incentives to set this price.

[6]Recall that there is the exogenous upper-bound for prices ζ, so $\overline{p} \geq \zeta$ is not considered.