JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 56, No. 6, Sept. 2021, pp. 2208–2236
© THE AUTHOR(S), 2020. PUBLISHED BY CAMBRIDGE UNIVERSITY PRESS ON BEHALF OF THE MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON
doi:10.1017/S0022109020000526

Income Fluctuations and Firm Choice

Scott R. Baker
Northwestern University
s-baker@kellogg.northwestern (corresponding author)

Brian Baugh
University of Nebraska
bbaugh2@unl.edu

Lorenz Kueng
Università della Svizzera Italiana, Swiss Finance Institute, Northwestern University
lorenz.kueng@usi

Abstract

How households shift spending across firms in response to income fluctuations is an important source of firm risk. Using transaction-level data, we study how households interact with the universe of retailers following income changes. We find that income increases within and across households result in substitution toward retailers in a category that are higher quality; smaller; more profitable; and have higher labor intensity, research and development (R&D) intensity, and equity betas. Although not all shifts are economically large, they do not average out across retailers. Thus, retailer choice has implications for key financial and macroeconomic outcomes, such as aggregate profitability and labor demand.

I. Introduction

Household consumption decisions exhibit substantial heterogeneity, both in the cross section of households and within households over time. Previous research has documented demand heterogeneity as a function of income and of household or product characteristics (e.g., Engle curves, choice of product variety). Moreover, heterogeneity in *how much* households respond to income changes (e.g., marginal propensities to consume (MPCs)), *when* they respond (e.g., anticipation effects, intertemporal substitution), and the *type of products* that are most responsive (e.g., durability, cross-product substitution) has also been shown to matter for any number of fields, including finance and macroeconomics.

One largely unexplored source of demand heterogeneity is the selection of firms by households, a key driver of heterogeneity in firm performance both cross-sectionally and over the business cycle. When selecting goods and services to consume, households make joint purchasing decisions across both products and firms. Whereas some firms offer unique products, others offer similar products to other firms, relying on marketing, customer service, or other amenities to persuade

We are very grateful to Jarrad Harford (the editor) for guidance and detailed comments and an anonymous referee for useful suggestions. We would like to thank seminar participants at Northwestern, the Hoover Institution, the International Institute of Public Finance (IIPF), Universidad Catolica de Chile, and the Copenhagen Workshop on New Consumption Data for helpful feedback.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

customers to patronize them. Decisions about *where* to purchase a good or service may shift concurrently with the type or quantity purchased. Although fluctuations in aggregate income (e.g., business cycles) are a key source of risk for firms, in general, the shifting of household spending *across* firms in response to income changes is an important additional source of risk for individual firms.

In this article, we focus on the type of firm that households directly interact with most frequently: retailers. Specifically, we analyze how income fluctuations influence how households shift spending across retailers. The paucity of research on retailer choice largely reflects the lack of retailer-specific data. Previous work offers only indirect and very limited evidence on how retailer choice varies across households and within households over time. Most research has been limited to using expenditure surveys or scanner data. Surveys typically record spending by disaggregated categories but do not provide information about the retailer from which households purchased these products (e.g., Consumer Expenditure Survey (CEX), Panel Study of Income Dynamics (PSID)). Scanner data, on the other hand, often come from a single retailer (e.g., Safeway) or contain only de-identified retailers (e.g., Nielsen), making it difficult to study heterogeneity in retailer choice linked to firm characteristics.

In this article, we use transaction-level data from a large personal finance online aggregator that allows us to comprehensively observe both household income and spending in detail. This class of data is increasingly used to study questions of household finance and empirical macroeconomics but has not been leveraged to investigate questions involving the firms that households interact with. We utilize these data to analyze how retailer choice within several sectors varies both across and within households.

In our data, each credit-card or debit-card transaction is associated with a textual description containing information that allows us to identify the retailer while protecting the anonymity of the household. For over half of household retail spending in our sample, we are also able to link retailers to external information from third parties such as Yelp, Orbis, Compustat, and the Center for Research in Security Prices (CRSP). Because the data span both private and public retailers, we are able to study heterogeneity by ownership type as well as other firm characteristics.

This article makes several contributions to the literature. First, we document new cross-sectional heterogeneities in consumer retailer choice, both within and across retailer categories. Second, we show how retailer choice responds to income changes within households over time. Third, we construct an index of retailer quality and show that households adjust the average quality of their retail spending quickly in response to income changes. Finally, we explore how these choices on the part of households interact with differences in firm characteristics and the extent to which they can drive aggregate financial and macroeconomic outcomes.

Not only do households shift their *level* of spending, but they also change where they shop. Cross-sectionally, we find that higher-income households frequent a larger number of unique retailers per month and that they tend to patronize smaller retailers. These patterns are not just a function of persistent household attributes (e.g., location or household type) but also hold within households over time as household income changes. We show that households are also more likely to

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

experiment with a retailer they have never visited before following an increase in income (e.g., trying out a new restaurant). This suggests that the "willingness to experiment" increases with income (both cross-sectionally and dynamically), consistent with "learning about demand" that prevents young firms from growing quickly in size (Foster, Haltiwanger, and Syverson (2016)).

Moreover, these shifts often occur along a retailer quality gradient. We find that the average quality of the retailers that they visit responds strongly to income patterns in both the short and long run. This result is most similar to that found by Coibion, Gorodnichenko, and Hong (2015), who use scanner data to show that consumers switch to lower-priced grocery stores during economic downturns. With our ability to directly observe the true retailer being patronized, we are able to extend this analysis by investigating whether these changes in retailer choice are correlated with firm characteristics. Importantly, if changes in retailer choice are correlated with firm characteristics, then changes in retailer patronage do not average out when aggregated across all households and can lead to real effects on firms and the economy.

For instance, we find that after an increase in income, households systematically shift their spending toward publicly traded retailers (affecting the volatility and riskiness of cash flow), to retailers with lower current profitability but higher research and development (R&D) investments, retailers with higher advertising expenditures, and retailers with higher measures of quality. We also find that households tend to shift toward retailers with higher betas following increases in income. Cross-sectionally, we find that retailers with higher-income customers have higher betas than those with lower-income customers. These effects are not *all* economically large in the short run. Households have substantial amounts of persistence in their choice of preferred retailers, but we demonstrate that these preferences do shift over time in nonneutral ways.

These demand-side findings complement a recent literature that uses plant-level U.S. Census data to document significant differences in the volatility of employment growth between private and public firms (e.g., Davis, Haltiwanger, Jarmin, and Miranda (2006), Davis and Kahn (2008)) and a similar literature in finance that studies recent changes in firm-level volatility (e.g., Campbell, Lettau, Malkiel, and Xu (2001), Fama and French (2004)).

The effect on key macroeconomic outcomes, such as employment, is more nuanced. Likely as a result of shifts along the aforementioned quality gradient, we find that an income decrease is associated with households shifting to less labor-intense retailers *within* a specific sector (i.e., "trading down") (Jaimovich, Rebelo, and Wong (2019)). However, households also tend to substitute from less to more labor-intense retail sectors following a decline in income. Because cross-sector substitution dominates within-sector substitution, the net effect is an increase in retail labor intensity.[1]

Our article relates to several strands of literature in finance and macroeconomics. Recent research in finance studies how firm heterogeneity and corporate decision making affect expected returns. A large and influential literature surveyed

[1]The net effect on total employment in the economy, of course, depends not only on the labor intensity of retailers but also on the labor intensity of the products purchased at those retailers (i.e., the entire supply chain).

by Fama and French (1992), (2008) explores factors that predict return differences across firms, and most of these studies have focused on supply-side explanations.[2]

Our article proposes a new mechanism to explain firm heterogeneity: differences in the customer base or "customer capital" and in the elasticity of consumer choices across retailers.

Our study therefore contributes to a smaller but complementary finance literature that relates *demand-side* differences to firm heterogeneity, both to intermediate outcomes, such as endogenous firm characteristics (e.g., capital structure, organizational form), and to firm performance (e.g., profitability, excess returns).[3]

With growing access to new micro-level transaction data, macroeconomists increasingly move beyond studying macroeconomic aggregates and instead explore the macroeconomic implications of various forms of consumer- and firm-level heterogeneities.[4] Our article connects heterogeneity in consumer demand with supply-side heterogeneities across firms.[5] For instance, our finding that consumers' first-time patronage of stores is positively correlated with income supports customer-base explanations of firm heterogeneity (e.g., price stickiness and firm size).[6]

Finally, our article belongs to a rapidly growing literature that uses transaction-level data to study firm behavior and markets (see, e.g., Einav, Klenow, Klopack, Levin, Levin, and Best (2018) and the literature cited therein).

The article is organized as follows: In Section II, we describe our data and measures of retailer characteristics. Section III documents cross-sectional facts and dynamic responses of retailer choices to income shocks at the household level. Section IV demonstrates some financial and macroeconomic impacts of retailer choice. Section V concludes.

[2]A nonexhaustive list of recently studied supply-side factors for explaining firm-performance heterogeneity includes organizational capital, management skill, and management practices (Adams, Almeida, and Ferreira (2009), Eisfeldt and Papanikolaou (2013)); price rigidity and capital structure (Gorodnichenko and Weber (2016), D'Acunto, Liu, Pflueger, and Weber (2018)); product market competition and production networks (Herskovic, Kelly, Lustig, and Van Nieuwerburgh (2017), Bustamante and Donangelo (2017), and Herskovic (2018)); and many more.

[3]Demand-side factors studied in previous research that explain excess returns include luxuries versus basic goods (Ait-Sahalia, Parker, and Yogo (2004)), durable versus nondurable goods (Gomes, Kogan, and Yogo (2009)), good- specific habit formation (Van Binsbergen (2016)), and shifts in demand elasticity (Stroebel and Vavra (2019), Nevo and Wong (2018)).

[4]Our cross-sectional facts are related to empirical studies of Engle curves, household budget shares, and product choice (variety and quantity–quality choice) (see Deaton (1997), (2016) for surveys). Our dynamic within-household results relate to a large literature on intertemporal consumption choice (e.g., MPCs, intertemporal substitution, home production, and other adjustment margins) (see Jappelli and Pistaferri (2017) for a recent survey).

[5]A recent literature explores demand-side explanations of differences in firm size. For instance, Foster et al. (2016) focus on homogeneous commodity-like product industries (e.g., ready-mixed concrete, manufactured ice) to control for quality difference and rule out other supply-side factors (productivity and cost differences). They conclude that demand-side factors must be the main explanation for persistent firm-size differences, particularly a learning-about-demand mechanism. Our article complements this strand of literature by providing direct evidence of this mechanism.

[6]Whereas most macroeconomic research on price stickiness focuses on supply-side explanations (e.g., menu costs, information costs), a few studies explore demand-side factors, particularly sticky customer capital (Blinder, Canetti, Lebow, and Rudd (1998), Rotemberg (2005), Hall (2008), and Dupraz (2016)). Our article contributes to this literature by using household-level data and variation in household income to document this channel. We study how a retailer's customer base changes with economic conditions.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

II. Data

A. Transaction-Level Linked-Account Data

Online aggregation of financial accounts is a popular service that allows households to easily monitor financial activities from across multiple financial institutions using a single web page or smartphone app. Account-aggregation services often allow features such as budgeting, expense tracking, and so forth. Dozens of companies currently provide such services, and our data come from one of the largest of these firms.

Once users initially sign up for the free service, they are given the opportunity to provide the service with usernames and passwords to a variety of financial accounts (checking, savings, credit card, brokerage, retirement, mortgage, student loan, etc.) from any financial institution, although our particular data are limited to bank and credit-card accounts. After signing up, the service automatically and regularly pulls data from users' financial institutions. The data contain transaction-level data similar to those typically found on monthly bank or credit-card statements, containing the amount, date, and description of each transaction. The full data set contains 2.7 million households from 2010 to 2015, and although the sample grows over time, there is very little attrition in our sample.

Because unobserved consumer spending from unlinked credit cards is a potential issue with this data set, we remove any household that makes excessive credit-card payments from the bank account relative to observed spending in the credit-card account. Specifically, we remove from the sample any household that, over our entire sample period, spends twice as much on credit-card payments than observed credit-card spending. This has the effect of removing households that we believe have substantial amounts of spending that we do not observe transactions for. A similar restriction could be made for regular transfers from unlinked checking accounts, although these are comparatively rare; American households tend to have a range of credit cards but generally only one or two checking accounts.

The sample is not a random sample of the population, but it appears to be widely representative, with some exceptions. Figure A.1 in the Supplementary Material illustrates the income distribution of our final sample relative to the U.S. Census. Although the raw sample differs from the true income distribution in the United States, the sample covers households with a wide range of incomes rather than solely identifying households of a particular income group. We also find that the households in our sample are well dispersed geographically, although we have higher concentrations of households in the states of California, New York, and Texas relative to true population distributions. However, dropping members from any given state (e.g., overrepresented states) or applying other weighting strategies does not substantially impact our results. Similarly, excluding households in the top or bottom deciles of income has little impact on our empirical results.

Recent work has also utilized similar transaction-based sources to make inferences about the financial habits of the broader population. For instance, Baker (2018), Baker and Yannelis (2017), and Kueng (2018) also utilize data from an online personal finance platform. They perform a multitude of validity tests comparing to data sources such as Census Retail Sales, home price data from Zillow, the Survey of Consumer Finance, and the CEX. They find a close parallel between

household-level financial behaviors and distributions in these sources relative to those found among users of the online platform. That is, conditional on basic demographic types, selection into the online platform did not predict differential financial behavior or characteristics.

Ganong and Noel (2019) and Olafsson and Pagel (2018) perform similar exercises using data taken from JPMorgan Chase and a financial services app covering the population of Iceland, respectively. Across a range of financial indicators, they find strong evidence of the external validity of their results using their sample population. Such results point to the fact that although these types of bank-derived sources will mechanically exclude financial activity by the unbanked, transaction-level financial data can produce accurate portrayals of aggregate economic activity and household behavior.

B. Retailer Identification

We choose to focus on household decision making in the context of retail spending because it offers a setting in which large numbers of potential providers are competing for a household's business. Relative to household spending on medical care, housing, education, or other services, there are generally many retailers within a given category (e.g., clothing, groceries) in a single geographic location that span a wide range of bundles of goods and quality of service. Moreover, because these firms interface directly with households through financial transactions, our data give us the clearest picture of this set of firms.

To perform our analysis, we must aggregate individual textual descriptions of transactions to particular retailers. Unfortunately, we cannot take each unique text description of a transaction as a separate retailer. This stems primarily from variation from two sources: i) variation in how different credit or debit cards report or abbreviate textual descriptions from firms and ii) variation within retailers across locations in the textual description accompanying a given transaction. That is, the textual description from a given store may appear differently when using a credit card rather than a debit card, and two locations of the same retail chain may have different textual descriptions, as well.

In light of these challenges in mapping transaction descriptions to firms, we approach this matching in two stages. First, we clean each transaction description string to remove common text, such as "Inc" and "Corp," and various other punctuation marks. We also remove long numerical series that represent a "transaction ID" unique to that household-retailer transaction.

This initial cleaning procedure reflects the sheer number of unique textual transaction descriptions in our database. Although some retailers' transaction descriptions are consistent across the sample, many retailers have a large number of individual transaction descriptions associated with them, especially those retailers that are highly geographically dispersed. For instance, the grocery store Safeway has 713 unique descriptions associated with it, even after stripping out generic strings and numeric transaction ID numbers. This is because most individual Safeway locations typically have a unique string description that includes both "safeway store" and the location of the store or some store identifier (e.g., "safeway store fairfax" or "safeway store canyon way").

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

The second stage of matching takes our cleaned set of unique transaction descriptions and works to map them to individual firms on a many-to-one basis. We start by conducting an algorithmic matching process (based on a Levenshtein distance operator) that assigns similarity scores to each transaction description in our sample relative to each retailer. This set of potentially matched retailers is taken as the entirety of retailers that are categorized as being in the retail industry by Compustat, as well as a large set of private retailers. Large private retailers are identified from a range of sources that aim to identify the largest nonpublic retailers according to annual sales, number of stores, or estimated value.

We supplement the automated approach with some hand-matching of the largest private and publicly traded retailers to ensure thorough coverage. In particular, we make sure that we hand-match the top 2,500 unique transaction descriptions (although not all of these 2,500 end up matched because many are nonretail transactions) to retailers to ensure that we are getting as clean of a view of retailer spending as possible among this sample.

The types of transactions that hand-matching typically identified were those in which a description included an abbreviation of a retailer that differed substantially from the full retailer name (e.g., "TGT" rather than "TARGET") or one in which there may have been a change in a parent firm–subsidiary firm relationship over time. In our entire sample of matched retailers, the mean number of unique text descriptions associated with a given retailer is 157, and the median number is 34.

Following this matching procedure, we are able to identify spending at specific retailers across households and across time and can also examine characteristics of retail spending within particular categories. In general, we aggregate retailer- or categorical-level data to a household-month level, depending on the particular specification being utilized. Each retailer is categorized into one of several categories of retail, such as groceries, restaurants and dining, clothing and shoes, specialty retail, general merchandise, gasoline and fuel, personal care, and home improvement. For the purposes of most of our analysis, we focus on the four largest categories of retailers: groceries, restaurants and dining, clothing and shoes, and general merchandise.

Overall, we believe that the total spending we observe approximates the total retail spending done by these households, given the completeness of the data and our sampling restrictions. Overall, the retailers we are able to match constitute approximately 57% of total retail spending. The fraction is highest for general merchandise (68%) and lowest for restaurants (24%). We are able to match 59% of clothing expenditures and 53% of grocery expenditures. Intuitively, this makes sense because restaurants tend to have a longer tail of single-location retailers, whereas general merchandise spending is more highly concentrated in big-box stores.

1. Retailer Financial Data

For results related to understanding how retailer substitution within households is related to firm characteristics reflected by firm financial statements, we limit our analysis to firms that we are able to match to Compustat and Orbis. Unfortunately, this can only be conducted for the subset of public retailers in our sample that we've matched to using the procedure outlined previously.

We define a retailer's labor intensity as the ratio of the number of employees a retailer has to its sales, in millions of dollars. We take an average of this value for a given retailer, using data from 2010–2015, where available. This approach yields a retailer-level measure of labor intensity.

To construct a time-varying household-level measure of average retailer labor intensity, we take an approach similar to that which we utilized to create the retailer quality variable in the previous section. Again, this is computed at a household-month level for our four retailer categories (general merchandise, groceries, restaurants/dining, and clothing/shoes) as well as an aggregate measure across retailers in all four categories.

In addition, we employ Compustat measures of annual profitability, advertising, R&D expenditures, and other financial attributes by retailer. Most of these measures are normalized at the firm level by total assets, and data are obtained for all firms with nonmissing values across our sample years, 2010–2015.

C. Measuring Retailer "Quality"

As one method of measuring how households systematically shift between different retailers over time, we attempt to construct a data-driven measure of retailer "quality." We proceed under the assumption that in the cross-section, higher-quality retailers will have a more affluent customer base. We test this assumption using external estimates of retailer prices (see Bils and Klenow (2001) for a discussion of household income, prices, and product quality). Retailer "quality" is then measured as the average household income of a store's patrons.

1. Comparing Customer Bases Across Retailers

We start from the perspective of individual retailers, seeking to measure the extent to which the income distribution of a retailer's customers differs from the income distribution in the cross section of households. We first calculate both the annual income of each household across all months that they remain in our sample and the overall household income distribution of the sample.

If all households shopped at all retailers in equal proportions, the fraction of revenue retailers would derive from households in a given income bin would simply be equal to the fraction of total spending made up by that income bin in the economy. For instance, if households that make \$50,000–\$51,000 constitute 1.5% of total retail spending, a retailer would be expected to receive 1.5% of its revenue from this group of households.

Figure 1 plots the difference between the observed distribution of retailer income and the "neutral" or "expected" distribution for a number of retailers. Households are binned by \$1,000 of annual income. Each graph plots the over- or underrepresentation of revenue from households for a selected sample of retailers at each point along the income distribution relative to the cross-sectional income distribution (i.e., the "excess mass" of the customer base income distribution for retailers of the same category). For instance, a value of positive 0.01 means that the named retailer received 1 percentage point more revenue from that portion of the income distribution than would be expected if all households shopped at all retailers with equal likelihood.

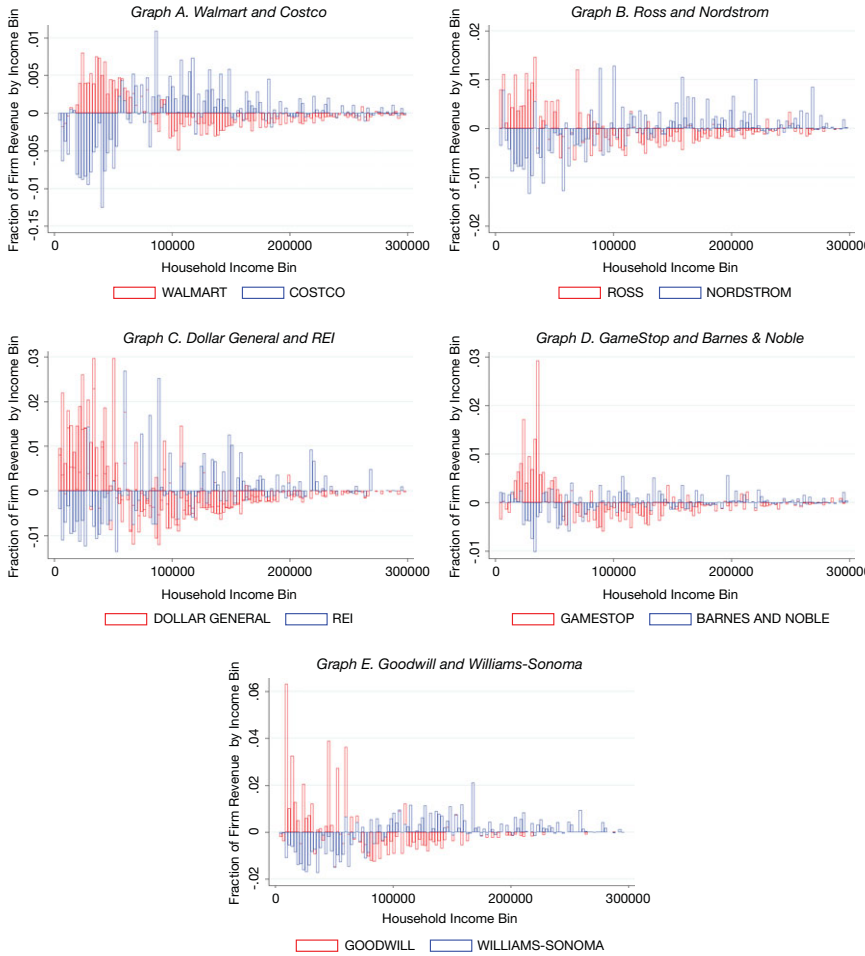
https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

We find large differences in the income distribution of customers between retailers within the same retailer category. For instance, in Graph A, we find that Walmart has a large relative surplus of lower-income customers, whereas Costco sees a deficit of such households and a surplus of higher-income customers. Similarly, Jack in the Box, Panda Express, Waffle House, Dollar General, and Ross tend to cater to lower-income households relative to their competitors: Panera Bread, P. F. Chang's, Corner Bakery, REI, and Nordstrom.

FIGURE 1

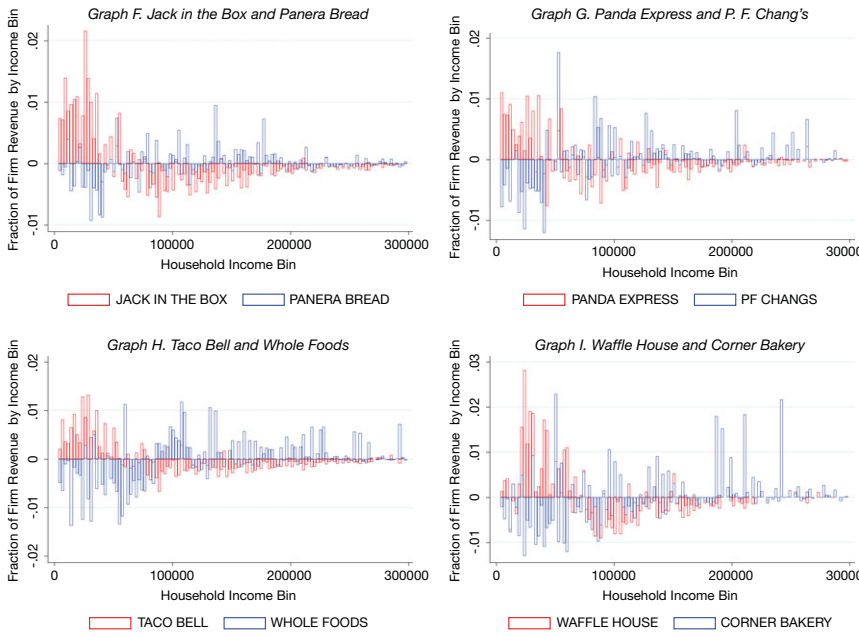
Income Heterogeneity in Retailers' Customer Base

The graphs in Figure 1 show the over- or underrepresentation of retailer revenue from households at each point along the income distribution for a selected sample of firms. First, the overall household income distribution of the sample is calculated. For each firm, the difference in the revenue of households from a particular point of the income distribution is calculated relative to the overall distribution and plotted in $2,500 bins. For instance, a value of positive 0.01 means that the named retailer received 1 percentage point more revenue from that portion of the income distribution than would be expected if all households shopped at all retailers with equal likelihood. Household income spans $0 to $300,000. Retailers are a selected sample of large retailers.



https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

FIGURE 1 (continued)



2. Households' Average Retailer Quality

We construct our measure of retailer quality by computing the average income, weighted by spending, of the retailers' customers. Using this method, higher-end firms will attract richer households, which will be reflected in a higher quality rating. This is consistent with the tendency noted by Bils and Klenow (2001), where the authors find that higher-income households tend to purchase higher-priced and higher-quality goods. One nice feature of using this measure of firm quality is that it is easily computed for all firms in our database, whether the firms are publicly traded or private and without regard to whether they can be linked to an external data source such as Compustat.

An important caveat to the utility of this measure of quality is that it is mechanically related to the average level of household income in the cross section. That is, because we assert that retailers of higher quality have customer bases with higher average income, households with higher levels of income will mechanically shop at higher-quality retailers in our data.

Although this self-referential nature limits the utility and meaningfulness of this quality measure in the cross section, variation in income *within* a household does not mechanically drive the quality of retailers that a household patronizes. That is, as a given household's income increases, it is not assured that the firm quality of the retailers the household frequents would increase (because each household makes up only a small portion of the customer base of any given firm). In essence, the shifting of the average quality of the stores at which the households shop informs us to what extent households facing income changes converge to the average type of store patronized by households of their new income level relative to their old income level.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

To determine the average quality of a household's portfolio of retailers for a given month, we proceed by calculating the dollar-weighted average of these retailer-level quality measures. For example, if a household spent $200 at Walmart (with quality computed as $67,792) and $400 at Target (with quality computed as $84,460) in a given month, the household's average quality in that month would be computed as follows:

$$\frac{\$200}{(\$200+\$400)}\times\$67,792+\frac{\$400}{(\$200+\$400)}\times\$84,460=\$78,904.$$

We compute measures of quality at the household-month level for our four retailer categories (general merchandise, groceries, restaurants, and clothing/shoes) as well as an aggregate measure across retailers in all four categories. Summary statistics for these variables are presented in Table 1.

We feel that this measure of firm-level quality corresponds well to a more subjective assessment of the relative distribution of quality among retailers in the United States. For example, the computed quality measures of selected general merchandise retailers are as follows: Walmart, $67,792; Target, $84,460; Barnes & Noble, $87,686; and Bed Bath and Beyond, $92,710. The computed quality measures of selected grocery retailers are as follows: Kroger, $78,434; Safeway, $87,588; Publix, $81,111; Whole Foods, $107,377; and Trader Joe's, $90,871. The computed quality measures of selected restaurants/dining retailers are

TABLE 1

Summary Statistics

Table 1 shows basic summary statistics for the households in our sample. The unit of measure is the household-month. Both unconditional values and values conditional on nonzero values are shown. MONTHLY_INCOME is the observed total household income for the household. MONTHLY_CREDIT_CARD_PAYMENTS is the sum of payments to credit-card companies. Spending data are limited to spending identified as retail spending (rather than, e.g., bills or mortgage payments or various services). Firm "quality" is determined by the dollar-weighted average household income of customers at a given retailer. Labor-intensity data refer to employees per million dollars of sales and are restricted to retailers in our data that are able to be matched to the Compustat database.

	Conditional		Unconditional	
	Mean	Std. Dev.	Mean	Std. Dev.
Panel A. Household Variables				
MONTHLY_INCOME	$5,114	$4,616	$3,729	$4,550
MONTHLY_CREDIT_CARD_PAYMENTS	$1,753	$2,640	$1,208	$2,337
MONTHLY_CREDIT_CARD_PURCHASES	$2,192	$2,899	$998	$2,240
HAS_CREDIT_CARD?	78%	42%	78%	42%
Panel B. Monthly Spending				
AGGREGATE	$764	$729	$759	$730
CLOTHING_&_SHOES	$203	$269	$109	$221
GENERAL_MERCHANDISE	$367	$413	$317	$404
GROCERIES	$310	$342	$240	$328
RESTAURANTS	$149	$151	$127	$149
Panel C. Annual Quality				
AGGREGATE	$78,417	$10,712	$78,417	$10,712
CLOTHING_&_SHOES	$82,415	$14,457	$82,415	$14,457
GENERAL_MERCHANDISE	$76,970	$9,776	$76,970	$9,776
GROCERIES	$81,203	$12,636	$81,203	$12,636
RESTAURANTS	$73,947	$11,818	$73,947	$11,818
Panel D. Labor Intensity (employees per $1M sales)				
AGGREGATE	0.00712	0.00349	0.00712	0.00349
CLOTHING_&_SHOES	0.00705	0.00176	0.00705	0.00176
GENERAL_MERCHANDISE	0.00486	0.00143	0.00486	0.00143
GROCERIES	0.00413	0.00167	0.00413	0.00167
RESTAURANTS	0.01609	0.00377	0.01609	0.00377

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

as follows: McDonald's, $66,651; Starbucks, $83,636; Chick-fil-A, $78,740; Chipotle, $79,359; Denny's, $62,673; and P. F. Chang's, $90,315. The computed quality measures of selected clothing/shoes retailers are as follows: Macy's, $89,163; Kohl's, $84,593; JC Penney, $73,721; and Nordstrom, $111,415.

This algorithm yields a continuous and quantitative measure of retailer quality for all retailers in the sample but is constructed without reference to any external benchmark. To validate our data-driven measure with a more "objective" metric of quality, we merge in information from the online review website Yelp.com for each retailer in our sample. Yelp.com publishes information regarding the price range that each retailer sets and denotes these prices by a number of one to four dollar signs ($–$$$$). Table A.1 in the Supplementary Material shows a significant positive relationship between our measure of retailer quality and the Yelp.com price level, providing external validity of our data-driven quality measure. This positive relation holds across all categories of retailers, with high-priced retailers having a significantly more affluent customer base on average.

III. Retailer Choice Elasticity Across and Within Households

To the best of our knowledge, this is the first comprehensive analysis of retailer choice using household-level transaction data.[7] We therefore start by examining the cross-sectional properties of households' choice of retailers. In particular, we seek to illustrate how retailer choice differs between households of different types along several dimensions. We then characterize changes in retail choice *within* a household as household income changes over time.

A. Comparing Retailer Choices Across Households

We first study differences in shopping behavior across the household income distribution, finding large differences in a number of dimensions. First, Graph A of Figure 2 plots the average number of retail transactions per month per household, the average number of unique retailers per month that a household visits, and the average monthly number of retailers that a household has never been observed patronizing in previous months (e.g., visits to a store that is "new" to that household). All three series show a positive and monotonic relationship with income. Households in the top decile of household income visit twice as many retailers as those in the bottom decile. This relationship tends to hold throughout each retailer category, with the higher-income quantiles visiting more unique retailers and new retailers for every category of retailer in our sample.[8]

Graph B of Figure 2 denotes the distribution of the dollar-weighted average retailer size for each quantile and for three categories of retailers: restaurants,

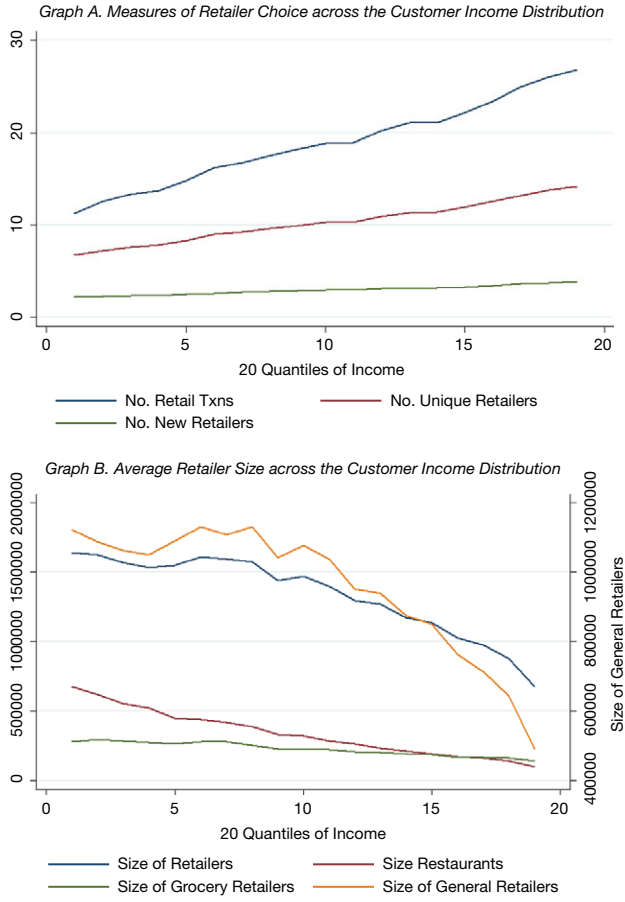
[7]There is a large literature in marketing and related fields that studies retailer and store attributes that determine retailer choice (see, e.g., Pan and Zinkhan (2006) for a survey), but we are not aware of other works that use customer- level financial data to study how households choose between retailers.

[8]We separately examine the overlap of retailers across households of different income bins and geographic areas. Although households of different incomes exhibit significantly different retailer choice, on average, there remains a substantial degree of overlap in the retailers that they patronize, especially for households living in similar geographical areas.

FIGURE 2

Retailer Choice Across Income Distribution

Graphs A and B of Figure 2 denote average values of a particular variable across 20 quantiles of observed household income. Graph A plots three series: the average number of retail transactions per month per household, the average number of unique retailers per month that a household visits, and the average monthly number of retailers that a household has never been observed patronizing in previous months. Graph B denotes the distribution of dollar-weighted average firm size for each quantile (left axis) and for three categories of retailers (right axis): restaurants, grocery stores, and general merchandise. Firm size for a particular firm is measured as the dollar-weighted number of transactions conducted at that firm across all individuals and all times in our sample. Thus, a firm with more outlets and more sales will be rated as a larger firm.



grocery stores, and general merchandise. "Size" for a particular retailer is measured as the dollar-weighted number of transactions conducted at that retailer across all individuals and all months in our sample. Thus, a retailer with more outlets and more sales will be classified as a larger firm. Here, we find that although the average size of retailers varies across the income distribution, the effect is markedly different for different types of retailers.

For instance, the average size of restaurants visited by the top decile of income is over 40% smaller than the average size of the restaurants visited by the bottom decile of income. Meanwhile, the average size of grocery stores hardly varies across the income distribution. The average size of general merchandise retailers

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

exhibits a hump-shaped relationship, with both the top and bottom deciles of the income distribution patronizing smaller stores than the middle of the income distribution.

This negative relationship between customer income and average retailer size provides novel *direct* evidence based, on household-level data, of the "learning about demand" channel proposed by Foster et al. (2016) to explain the wide size distribution of firms in the United States. After ruling out supply-side explanations using plant-level U.S. Census data, the authors conclude that "these patterns do not reflect productivity gaps, but rather show differences in demand-side fundamentals," such as building a customer base. Much of the declines in firm size among the higher-income households tend to reflect movement away from big-box retailers (e.g., Walmart and Target) and large restaurant chains (e.g., McDonald's and Chili's) and toward smaller and more local firms.

We also explore the extensive margin of household retailer choice: the frequency with which households visit retailers that they have never previously visited. Figure 3 shows the distribution of spending at "new" retailers across several categories as a function of household income. For each household-month, we calculate the fraction of spending done within a particular category at a retailer that the household has not previously shopped at.

A value of 1 means that all of the household's shopping in that retailer category was done at a retailer the household has not previously visited, whereas a value of 0 means that all of a household's spending of that type was at retailers the household has shopped at before. The first 12 months of a given household are excluded to "burn in" previously visited retailers. Each graph in the top row shows a histogram of these values across all household-months in our sample. The bottom row mirrors the top row but excludes values of 0 and 1 to provide more detail.

We find that trying out a new retailer for only a portion of a household's monthly spending is rare; spending at new retailers is concentrated around 0% or 100% of monthly spending within a category.[9] However, household spending at new restaurants is much less concentrated at the tails of the distribution relative to spending at general merchandise retailers and grocery stores. That is, households do a much larger proportion of restaurant spending at places they've never visited before. In contrast, monthly spending at general merchandise and grocery retailers is more commonly done entirely at retailers that a household has patronized in previous months (yielding a value of 0) or entirely shifted to a new retailer (yielding a value of 1).

B. Comparing Retailer Choice Within Households

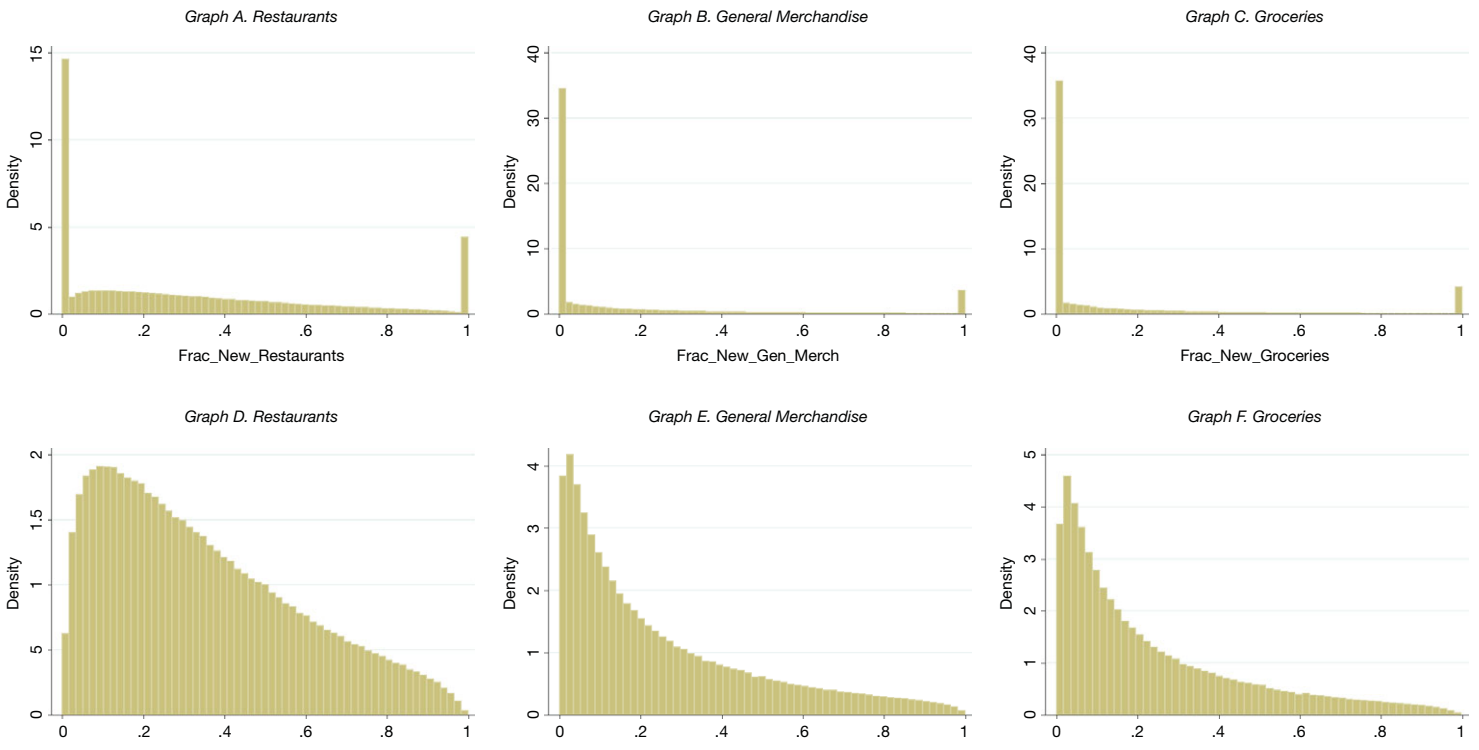
We now turn to examining how households change their choice of retailers over time following fluctuations in household income. Each panel of Table 2 shows

[9]Much of the concentration at the tails is driven by months in which households only have one or two transactions in a particular category (e.g., going out to eat only twice or visiting the same grocery store a few times). Restricting this analysis to months in which there is a minimum threshold of spending in a given category reduces the tail outcomes, 0% or 100% of spending at a new retailer, while leaving the interior distributions relatively unchanged.

FIGURE 3

Fraction of "New" Spending (Inclusive and Exclusive of 0 and 1)

In Figure 3, for each household-month, we calculate the fraction of spending done within a particular category at a retailer that the household has not previously shopped at. A value of 1 means that all of the household's shopping in that retailer category was done at a retailer the household has not previously visited, and a value of 0 means that all of a household's spending of that type was at retailers where the household has shopped before. The first 12 months of a given household are excluded to "burn in" previously visited retailers. Graphs A–C show a histogram of these values across all household-months in our sample. Graphs D–F mirror Graphs A–C but exclude values of 0 and 1 to show more detail.



https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

TABLE 2

Retailer Choice Elasticity Using Customer-Level Income Changes

In Table 2, the dependent variable for each specification varies by panel. Each column varies the set of retailers that the given variable spans (i.e., aggregated across retailers (AGG), solely restaurants (REST), solely grocery stores (GROCERIES), solely general merchandise (GEN_MERCH), solely clothing and shoes (CLOTHES)). Each specification includes household and period fixed effects (FE) as well as logged household income. Panel A considers logged spending, Panel B considers the logged number of unique retailers, Panel C's dependent variable is the Herfindahl–Hirschman index (HHI) within the category of retailer, Panel D looks at the weighted average of store size that a household shops at (as measured by logged number of total transactions across our sample at a given retailer), and Panel E measures the weighted average of firm quality for a household-month. Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	AGG 1	REST 2	GROCERIES 3	GEN_MERCH 4	CLOTHES 5
Panel A. Aggregate Spending					
ln(INCOME)	0.230***	0.180***	0.139***	0.215***	0.130***
	(0.00157)	(0.00180)	(0.00219)	(0.00220)	(0.00309)
R^2	0.516	0.487	0.513	0.435	0.278
Panel B. Unique Stores					
ln(INCOME)	0.103***	0.0903***	0.0365***	0.0600***	0.0326***
	(0.000939)	(0.000995)	(0.000806)	(0.000847)	(0.00103)
R^2	0.592	0.582	0.448	0.443	0.258
Panel C. Spending Concentration Across Retailers (HHI)					
ln(INCOME)	−0.0660***	−0.0858***	−0.0308***	−0.0515***	−0.0392***
	(0.00104)	(0.00113)	(0.000914)	(0.000963)	(0.00129)
R^2	0.463	0.545	0.405	0.369	0.270
Panel D. Average Store Size					
ln(INCOME)	−0.0176***	−0.0624***	−0.0418***	0.0387***	−0.0251***
	(0.00291)	(0.00430)	(0.00528)	(0.00497)	(0.00638)
R^2	0.494	0.526	0.472	0.409	0.294
Panel E. Average Firm Quality					
ln(INCOME)	0.00411***	0.00521***	0.00353***	0.00273***	0.00269***
	(0.000171)	(0.000215)	(0.000263)	(0.000238)	(0.000403)
R^2	0.719	0.635	0.646 42	0.594	0.476
No. of obs.	1.300.499	1.150.360	1.036.070	1.118.010	676.405

income elasticities ϵ_x of different outcome variables x_{it} after controlling for period (τ_t) and household fixed effects (α_i),

$$\ln(x)_{it} = \epsilon_x \ln(\text{income})_{it} + \alpha_i + \tau_t + u_{it} \quad (1)$$

Important to note is that this specification solely identifies short-run (concurrent) impacts of income fluctuations on retailer choice. This likely mutes the size of the observed response relative to any longer-run adjustments in retailer choice that the household engages in. In Table 3, we show that the long-run responses tend to be substantially larger than the short-run effects. This gradual adjustment and switching across retailers makes sense because households may take time to experiment and gather information about new retailers following an income shock. In the extreme, in response to a change in permanent income, households may relocate in the long run to somewhere with a different set of local retailers that is generally of higher or lower quality than their previous location.

Column 1 of Panel A in Table 2 provides the overall elasticity across all retailers, and columns 2–5 break elasticities down by retailer category (general merchandise, groceries, restaurants, and clothing/shoes).

This panel approach produces similar qualitative patterns as the cross-sectional analysis in the previous section. However, the magnitudes of the effects are substantially reduced within households relative to those measured across

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

TABLE 3

Dynamic Responses to Income Changes

Logged income is measured at a household-month level and incorporates all observable income for the household. Three leads of logged income and 6 lags of logged income are included in Table 3, alongside contemporaneous logged income. SPENDING refers to household logged spending across all retail categories. UNIQUES refers to unique retailers that a household visits in a given month across all retail categories. HHI refers to the Herfindahl–Hirschman index of dollars of retail spending for a household across all retailers and is calculated by squaring the fraction of total spending of the household done at each retailer and summing across all retailers. A lower HHI indicates a greater amount of dispersion in spending across multiple retailers. Each specification includes household and period fixed effects (FE). Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	SPENDING 1	UNIQUES 2	HHI 3
ln(INCOME)_3_MONTH_LEAD	0.00555**	0.00374***	−0.00492***
	(0.00233)	(0.00142)	(0.00161)
ln(INCOME)_2_MONTH_LEAD	0.0142***	0.00708***	−0.00433**
	(0.00245)	(0.00149)	(0.00169)
ln(INCOME)_1_MONTH_LEAD	0.0223***	0.0125***	−0.00934***
	(0.00248)	(0.00150)	(0.00171)
ln(INCOME)	0.148***	0.0625***	−0.0361***
	(0.00249)	(0.00151)	(0.00172)
ln(INCOME)_1_MONTH_LAG	0.101***	0.0448***	−0.0275***
	(0.00248)	(0.00151)	(0.00171)
ln(INCOME)_2_MONTH_LAG	0.0384***	0.0144***	−0.00701***
	(0.00248)	(0.00151)	(0.00171)
ln(INCOME)_3_MONTH_LAG	0.0111***	0.00282*	−0.00202
	(0.00248)	(0.00150)	(0.00171)
ln(INCOME)_4_MONTH_LAG	0.0147***	0.00440***	−0.00240
	(0.00246)	(0.00149)	(0.00169)
ln(INCOME)_5_MONTH_LAG	0.00538**	0.00193	−0.00180
	(0.00243)	(0.00147)	(0.00167)
ln(INCOME)_6_MONTH_LAG	0.00789***	0.00468***	−0.00251
	(0.00231)	(0.00140)	(0.00160)
No. of obs.	813,341	811,023	821,324
R^2	0.557	0.636	0.500
Year-month FE	Yes	Yes	Yes
Household FE	Yes	Yes	Yes

households. This is largely driven by two factors. First, these within-household results measure short-run elasticities, whereas households experiencing changes in long-run income converge only gradually to the retailer choices of a household with a different average income. Second, different geographic areas have both different average income and different choice sets of retailers, so a household may be unable to "take advantage" of a higher level of income by shopping at different retailers without moving its physical location.

Panel A of Table 2 looks at the relationship between household income and household retail spending at a monthly level. Overall, we find that the elasticity of monthly retail spending with respect to income is approximately 0.2. That is, for a 10% increase in household income, retail spending tends to increase by 2% in the same month. A strong relationship between income and retail spending is unsurprising. Both the sign and the magnitude are consistent with a large literature that estimates short-run MPCs out of income changes (see Jappelli and Pistaferri (2017) for a survey of relevant literature).

We note that retailer spending elasticities vary substantially across retailer type. Spending at general merchandise retailers tends to be the most responsive, whereas spending at grocery stores and clothing retailers is only about half as elastic within a household.

In Panel B of Table 2, we begin to examine other elements of household retailer choice in response to income changes. Mirroring our cross-sectional

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

approach, we test whether households tend to visit more unique retailers when household income increases. We find that this is the case across all categories, with the number of unique retailers increasing between 0.4% and 1% following a 10% increase in household income. This fact can also be expressed as a decline in the dollar-weighted Herfindahl–Hirschman index (HHI) of household spending. Panel C shows the HHI declining by approximately 0.7% for every 10% increase in household income. Both panels reflect an increasing desire for variety across a range of retailers as households become richer.

Panel D of Table 2 demonstrates that the average size of the retailers that households frequent (as measured by retailer sales or the number of transactions observed at that retailer) declines when household income increases. This reflects a decline in shopping at big-box stores and national fast-food chains and an increase in shopping at more local and boutique retailers. Households are more likely to spend their additional income at smaller retailers and at retailers that they have never frequented before. This supports the demand frictions proposed by Foster et al. (2016), who speculate that "new businesses are small because demand for their product is low, and demand is low because of informational, reputational, or other frictions."

A move away from larger stores may also be reflected in the changes observed in average retailer quality as household income increases. Consistent with this channel, Panel E of Table 2 shows that retailer quality is positively associated with income, even within a household over time.

We might worry that the short-run impacts of changes in income on retailer choice may be reversed in the longer run. In Table 3, we show that these effects on spending and retailer choice are not reversed; they tend to increase in magnitude over time. The dynamic specification adds 3 leads and 6 lags of household income to equation (1) to observe household responses in advance of and following changes in household income.

We find generally small impacts in advance of any change in income, suggesting that most of the variation in income that we are utilizing is relatively unanticipated by the household. However, we find large and significant impacts in the months following changes in household income and no evidence of reversals in the direction of the effect. This is consistent with changes in spending and retailer choice evolving over time, with impacts after 6 months being, on average, more than twice as large as the contemporaneous effects.

Overall, households that see changes in income tend to immediately adjust not only their levels of spending but also the composition of retailers that they visit. The largest effects are on the overall number of retailers that households visit. Firm quality and size also respond, but to a smaller extent, in the first months following the income change. These changes tend to be most pronounced among restaurants and general merchandise retailers.

It is worth reiterating that even though a cross-sectional analysis of household income and firm quality would produce a mechanical relationship, the previous *within*-household analysis is immune from such issues. Nonetheless, in the following section, we exploit additional sources of income variation within households to strengthen our results.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

1. Retailer Choice Following Job Loss and Tax Rebates

Our prior results hinge on changes in within-household levels of income identifying the relationship between income and various aspects of retailer choice, more generally. These fluctuations in income represent a combination of both income changes that are anticipated by households and those that are largely unanticipated. In addition, we primarily examine short-run responses to income changes, whereas some changes in income will be perceived to be more persistent than others.

In this section, we examine the response of retailer choice to two other sources of changes in household income that may stem from more exogenous and unexpected sources than the average within-household monthly fluctuation.

First, we test how retailer choice responds to an unemployment spell in Table A.2 of the Supplementary Material, which mirrors Table 2. We find here that unemployment spells reverse the effects seen from increases in household income. When faced with a period of unemployment, households cut spending by approximately 19% and decrease the number of retailers they visit by approximately 12%. They also increase their spending concentration across retailers and decrease the quality of the retailers they visit. Although most results largely coincide with our results utilizing overall monthly changes in household income, we do see that store size tends to *decrease* at the start of an unemployment spell. Some of this shift appears to be driven by households moving from big-box stores to lower-quality stores offering small purchase quantities (e.g., shifting from Target to Dollar General).

Our second test of an alternate source of household income variation comes from the arrival of tax refunds, as detailed in Table A.3 of the Supplementary Material. We look for responses of retailer choice that coincide with the arrival of a tax refund (we scale tax refunds by household income such that a value of 1 would represent a refund equal in size to a household's average monthly income). We again see results that mostly align with our primary specifications: Spending and retailer variety increase, average retailer size decreases, and firm quality increases.

Across both of these alternative specifications, we find that households seem to respond to these income shocks in ways similar to the typical monthly change in income in our main specifications. These results reinforce the likelihood that most shocks to household income will drive households not only to change their level of spending but also to systematically adjust the composition of retailers that they shop at.

C. Retailer Choice and Firm Characteristics

In addition to affecting levels of spending at retailers overall, the extent to which income fluctuations induce households to shift spending toward or away from retailers with particular financial characteristics has the potential to significantly shift average firm characteristics over the business cycle.

Thus far, our analysis has evaluated retailer choice across *all* retailers that the households in our sample visit. However, to analyze how these choices of retailers are related to the financial and labor characteristics of retailers, we must impose some restrictions on the data. To gain such information about retailers in our sample, we must restrict our analysis to transactions able to be matched to firms that can be identified in Compustat and Orbis.

TABLE 4

Elasticity of Firm Attributes in Response to Customer-Level Change in Income

In Table 4, the dependent variable for each specification varies by panel. Each column varies the set of retailers that the given variable spans (i.e., aggregated across retailers (AGG), solely restaurants (REST), solely grocery stores (GROCERIES), solely general merchandise (GEN_MERCH), solely clothing and shoes (CLOTHES)). Each specification includes household and period fixed effects (FE) as well as logged household income. Panel A covers both private and public firms, looking at the fraction of spending done at public firms. For all other panels, the dependent variable is an average of a particular firm financial characteristic, weighted by dollars of spending, across all public firms that the household conducts spending at in a given month. Panel B looks at the ratio of gross profits to total assets, Panel C looks at the ratio of research and development (R&D) spending to total assets, Panel D looks at the ratio of employees to revenue, and Panel E looks at the average beta of public firms, as measured by 5-year unlevered Dimson betas. Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	AGG 1	REST 2	GROCERIES 3	GEN_MERCH 4	CLOTHES 5
Panel A. Fraction Public					
ln(INCOME)	0.00505***	−0.0506***	−0.00830***	−0.0333***	−0.0259***
	(0.00173)	(0.000899)	(0.00117)	(0.000990)	(0.00101)
R^2	0.445	0.297	0.682	0.226	0.461
Panel B. Profitability					
ln(INCOME)	0.0282***	0.0192***	0.00861***	0.00122	−0.000796
	(0.000956)	(0.00125)	(0.00114)	(0.00101)	(0.00186)
R^2	0.288	0.272	0.656	0.282	0.279
Panel C. R&D Intensity					
ln(INCOME)	0.0334***	0.00960***	−0.00428***	0.0313***	0.00239
	(0.000933)	(0.00129)	(0.00163)	(0.000965)	(0.00206)
R^2	0.344	0.253	0.439	0.369	0.157
Panel D. Labor Intensity					
ln(INCOME)	−0.0298***	0.00603***	0.0194***	0.0197***	−0.00606***
	(0.000998)	(0.00102)	(0.00132)	(0.00104)	(0.00200)
R^2	0.294	0.270	0.673	0.336	0.346
Panel E. Betas					
ln(INCOME)	0.0143***	0.0168***	0.0106***	0.0263***	0.0198***
	(0.000875)	(0.00125)	(0.00115)	(0.000894)	(0.00182)
R^2	0.423	0.251	45 0.688	0.444	0.374

Table 4 examines the extent to which we can observe shifts in average firm characteristics by looking at changes in income within a household. The dependent variables in these regressions take the form of averages of the noted variables for a household-month, weighted by the amount of spending done at each retailer by that household (i.e., by the retailer-specific expenditure share of each household). For instance, the dollar-weighted average profitability of retail firms for household i in month t is given by the following:

$$\text{PROFITABILITY}_{it} = \sum_{f=1}^{F} \text{PROFITABILITY}_f \times \frac{\text{SPENDING}_{ift}}{\text{SPENDING}_{it}},$$

where PROFITABILITY_f is taken as the *average* level of profitability for retailer f for the entirety of our sample period, 2010–2015. Each dependent variable of Panels B–E in Table 4 is normalized such that the standard deviation is equal to 1.

In Panel A of Table 4, we find that households spend a larger share of their income at public retailers when their income is higher. Looking within individual categories of retailers, we find the reverse; each category sees a reduction in the proportion of spending at public retailers after an income increase. This is consistent with households patronizing smaller and more local (and hence, generally private)

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

retailers as their income increases. The increase that we find in aggregate spending is driven by a relative substitution toward the general merchandise category, where the fraction spent at public retailers is highest, and away from the grocery and restaurant categories, where the fraction of spending at public retailers is the lowest.

Although these effects are highly significant and robust across demographic groups, locations, and time, they are, for the most part, economically small shifts in choices across public and private retailers. A doubling of monthly income only increases the proportion of spending done at public retailers by approximately 1% relative to the average proportion. The largest effect that we find here is for restaurants, where a doubling of income drives a fall in spending at publicly traded restaurants of approximately 10% relative to the average proportion.

Panels B–D of Table 4 take a similar approach to the financial characteristics of the retailers that households patronize, measuring the dependent variable in terms of standard deviations across the sample. Because of the variables in question, we can calculate these measures solely for spending done at public retailers based on Compustat data. Here we see that as household income rises, households shift spending toward retailers with higher levels of R&D intensity, higher levels of labor intensity, and higher levels of profit. These trends seem to hold true both across categories of retailers as well as within categories, although shifts across grocery stores do not follow a similar pattern in terms of advertising and R&D.

In contrast to these quantitatively smaller results, in Table 2, we saw a substantial increase in unique store visits following fluctuations in income. The smaller magnitudes of response across these average firm characteristics demonstrate that while households shift spending substantially across individual firms, they may tend to stick to firms that are similar in various ways to the firms that they have "left behind," at least in the short run.

Allowing for longer-run responses through dynamic lag specifications, we find that these effects grow gradually over the 3–6 months following income changes. On average, we find that the aggregated effects in the two quarters following an income change are approximately 2.2 times larger than the contemporaneous firm-characteristic effects. Extending the specifications beyond this 3- to 6-month window generally produces imprecisely estimated coefficients.

Finally, Panel E of Table 4 shows how households shift spending across retailers in terms of the weighted average of unlevered betas. That is, does the average dollar of household spending go to a public retailer with a higher level of idiosyncratic volatility when income increases? We find that across all categories and also at an aggregate level, households experiencing increases in income shift spending toward retailers with higher betas.[10]

[10]We utilize 5-year unlevered "Dimson betas" to account for nonsynchronous trading in the market index relative to small individual equities. These betas are designed to limit biases for small stocks and for stocks in the early period of the sample. They are computed using the following regression:

$$r_{i,t} = a + b_1 r_{m,t} + b_2 r_{m,t-1} + r_{m,t-2} + \epsilon_{i,t}.$$

For each retailer, the regression is run at the end of every month, using the previous 5 years of daily return data (roughly 1,250 trading days). Finally, betas are constructed such that $\beta = b_1 + b_2 + b_3$ and are then unlevered. We take the average across all years in our sample, 2010–2015, for each retailer.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

TABLE 5

Firm Quality and Stock Market Betas

The dependent variables in Table 5 are the average betas of public firms across our sample period, as measured by 5-year unlevered Dimson betas. Firm quality is determined by the dollar-weighted average household income of customers at a given retailer and is averaged across our sample period. Also included is the logged average annual firm revenue taken from Compustat. Column 4 weights observations by firm revenue. Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	STOCK_BETA 1	STOCK_BETA 2	STOCK_BETA 3	STOCK_BETA 4
ln(FIRM_QUALITY)	0.546***	0.615***	0.520**	0.535***
	(0.185)	(0.182)	(0.202)	(0.201)
ln(REVENUE)		−0.0490**	−0.0556**	−0.0572**
		(0.0192)	(0.0268)	(0.0266)
No. of obs.	120	120	120	120
R^2	0.069	0.119	0.268	0.273
Category fixed effects	No	No	Yes	Yes

Given the propensity of individuals to shift toward higher-beta firms after income increases, we next evaluate whether cross-sectional differences in beta across retailers can be explained by the income profiles of shoppers. For low-quality retailers (i.e., those with lower-income customers), gains from rising aggregate income during periods of economic expansion will be partially offset by losses from substitution toward high-quality retailers. Conversely, losses from falling aggregate income during periods of economic contraction will be partially offset by gains from substitution toward low-quality retailers.

In Table 5, we directly test for this predicted positive correlation between beta and firm quality in the cross section among the publicly traded retailers in our sample. In column 1, we regress 5-year unlevered Dimson betas on our firm-quality measure. The results are both statistically and economically significant. A 1-standard-deviation increase in firm quality is associated with an increase in beta of 0.1 (approximately 10 percentage points). Further, the R^2 in column 1 indicates that our quality measure explains 6.9% of the variation in unlevered betas. In columns 2 and 3, we add controls for firm revenue and category fixed effects. In column 4, we weight observations by average firm revenue. The coefficient on firm quality remains statistically and economically significant throughout each specification. Overall, our findings provide additional granularity and context to the findings of Ait-Sahalia et al. (2004) and reinforce the importance of substitution across firms as an important source of risk to firms.

IV. Retailer Quality and Firm Choice

As in Section II, we construct measures of firm quality derived solely from our sample of transactional data as the average level of income of customers of a given retailer. One strength of this transaction-based approach is the ability to glean more information about private-sector firms that was previously unavailable (or only available on a case-by-case base for individual firms but not the entire retail sector). Although data on the characteristics of public firms are readily accessible (e.g., from Compustat or CRSP), details regarding the properties of a wide swath of private firms remain relatively opaque, and such a transaction-based approach can yield new insights across the universe of firms that households interact with.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

A. Quality Smoothing Within Month

Past work has noted that households may not fully smooth even *between* the receipt of periodic income. That is, households do not fully smooth consumption between paychecks, even when both the timing and the amounts of the paychecks are known in advance. Articles such as those by Shapiro (2005), Stephens (2003), and Mastrobuoni and Weinberg (2009) find evidence for such intra-month consumption cycles for a wide range of households.

In Table 2, we observed that households tend to adjust the average quality of the retailers they patronize in response to changes in household income over time. We follow in the vein of the previous literature by testing whether changes in retailer quality also follow an intra-month pattern consistent with households not fully smoothing between paychecks on this dimension.

We construct daily measures of average retailer quality for each household: the weighted average of firm quality given the amount spent at each store on a particular day. Using the online personal financial website's categorization for "Paychecks," we also map each day to the length of time since the last paycheck was received by that household.[11]

In Table 6, we see that the time since the last paycheck does have a negative and significant relationship with the quality of retailer that a given household is patronizing. For households on a monthly paycheck frequency, the average retailer quality just before receiving their next paycheck is approximately 1.5% lower than the average quality of retailer visited on payday itself. For restaurants, this effect is much steeper, with retailer quality among restaurants exhibiting more than twice the decline as for the average retailer.[12]

In columns 6 and 7 of Table 6, we expand the exercise from columns 1 and 2. Here we mirror our earlier specifications but include interactions with the average MPC out of income that the household exhibited over time (i.e., across months) in our sample. Households that have higher MPCs may also be those that exhibit the least amount of smoothing in terms of retailer quality in the short run. We find evidence that households with higher MPCs are the ones that tend to downgrade the quality of the retailers they visit most strongly as the time since their last paycheck increases.[13]

Overall, households in the highest decile of MPCs tend to decrease quality at 2 to 3 times the rate of the median household. Whereas other work has looked at levels of spending or time spent on home production as avenues to discuss high-

[11]We limit households in this sample to those that typically receive fewer than three paychecks per month. For households receiving more than three paychecks per month, it may be difficult to accurately ascertain any lack of smoothing because of the short length of time between paychecks. Indeed, if we limit our sample to those that receive many paychecks per month, we find no evidence of shifts in the quality of retailers between paychecks.

[12]We note that the observation of consumption, relative to spending, is difficult at such a high-frequency level: Most purchases are durable to at least some extent over a period of days or weeks. However, the consistent ordering of purchases over this period points to a possible effect on actual consumption. Also reassuring is the fact that we see similar results when restricting solely to restaurant spending, which is likely the least durable spending possible.

[13]We find negative impacts across all categories, although only retaining a statistically significant coefficient for aggregate retail spending and for restaurants.

TABLE 6

Individual MPC and Intra-Month Changes in Retailer Quality

In Table 6, MONTHS_SINCE_PAY denotes the time since the last receipt of income that is coded as a paycheck for a given individual. Regressions are run at a daily level. Each column's dependent variable is the logged weighted average of firm "quality" for a given type of retailer (i.e., all retailers (QUAL_ALL), restaurants (QUAL_RES), general merchandise (QUAL_GEN), groceries (QUAL_GRO), clothing and shoes (QUAL_CLOTH)). Each specification includes individual and period fixed effects (FE). MPC is a household's average marginal propensity to consume out of income at a monthly level, as calculated across the entire sample period (excluding the first 3 months of each household's data). Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	1 QUAL_ALL	2 QUAL_RES	3 QUAL_GEN	4 QUAL_GRO	5 QUAL_CLOTH	6 QUAL_ALL	7 QUAL_RES
MONTHS_SINCE_PAY	−0.0150***	−0.0315***	−0.00319**	−0.0100***	−0.00619*	0.00251	−0.0355***
	(0.000909)	(0.00139)	(0.00138)	(0.00171)	(0.00337)	(0.00330)	(0.00509)
MONTHS_SINCE_PAY * MPC						−0.0335*	−0.0854***
						(0.0188)	(0.0292)
No. of obs.	11,040,170	5,305,626	5,000,604	3,500,327	919,307	11,040,170	5,305,626
R^2	0.455	0.499	0.416	0.594	0.457	0.436	0.484
Period FE	YES	YES	YES	YES	YES	YES	YES
Individual FE	YES	YES	YES	YES	YES	YES	YES

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

frequency consumption smoothing, our work points to the fact that retailer quality is another avenue of consumption adjustment that households take advantage of as well. Such fluctuations can also have implications for firms. Based on our results, higher-quality firms tend to be the most exposed to both unexpected and expected reductions in income, and this can translate into higher revenue volatility in both the short and long run.

B. Macroeconomic Implications of Retailer Choice

Shifts among retailers of different quality can also have implications for the broader macroeconomy as a result of firm attributes that are correlated with this quality gradient. Previous work, such as that by Jaimovich et al. (2019), has noted that substitution across retailers in the face of fluctuations in household income may exacerbate swings in labor-market demand over the business cycle as a result of variation in the average labor intensity of different retailers.

Panel D of Table 4 directly examines this relationship: whether increases in household income lead to increases in the average labor intensity of the retailers that household visits. We find this to be the case for grocery stores, restaurants, and general merchandise retailers. However, we find a negative relationship when looking at the same regression for clothing retailers.[14]

Figure 4 demonstrates that quality is one channel through which this relationship may operate. We find that, on average, retailers of higher measured quality are more labor intensive in a given retailer category. This tends to be true across all categories in our sample, except for clothing. These results hold true in regression form as well. Across all matched retailers, we find that increases in firm quality tend to be negatively associated with a firm's labor intensity. This remains true when weighting by retailer size or by censoring the top or bottom 10% of the sample by labor intensity and when running the regression in logs or in levels.

There are two potential substitution channels at work. As consumers' income increases, they might substitute to new retailers with systematically different labor intensities (e.g., from lower- to higher-quality restaurants), and they might spend their additional income differently across retailer categories (e.g., more on restaurants and less on groceries). Because of significant differences in labor intensity across retailer category, the total effect on employment in the retail sector depends on both channels.

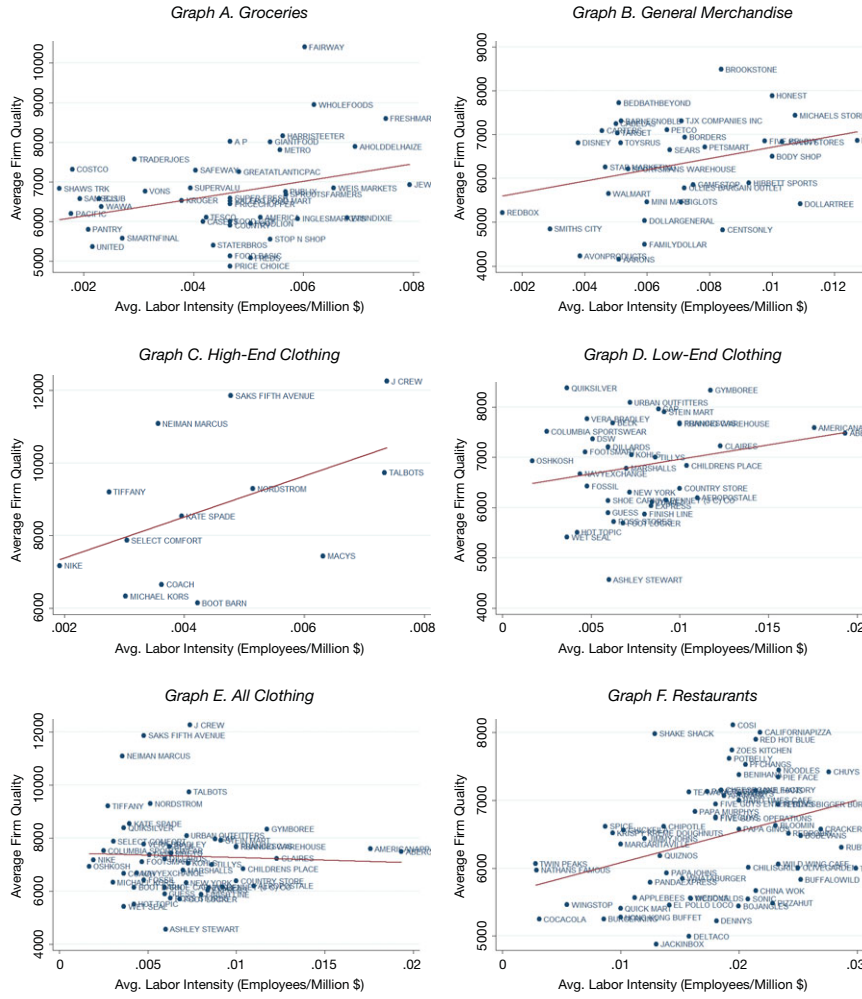
To interpret the average magnitudes of these effects, we imagine a scenario where household income doubles for a 12-month period (using coefficient estimates from our dynamic specification). At the end of this period, a household would be visiting restaurants that were, on average, 13% more labor intensive, general merchandise retailers that were 6% more labor intensive, grocery stores that were 0.5% more labor intensive, and clothing retailers that were 8% *less* labor intensive. These results are largely consistent with those noted by Jaimovich et al. (2019), who provide some aggregate evidence of households "trading down" to lower-quality and less labor-intensive firms during recessions.

[14]As in Table 3, we show the dynamics of the average labor intensity of the retailers a household visits in Table A.4 in the Supplementary Material. We see similar results, with effects generally building over the several months following the change in income.

FIGURE 4

Firm Quality and Labor Intensity

Each graph in Figure 4 shows all matched retailers in a given retailer category and plots their firm "quality" and the labor intensity of the firm. Quality is measured as the dollar-weighted average household income of a retailer's customers in our data. Labor intensity is measured as employees per million dollars in sales across all matched years in the data (2010–2015). A line of best fit is also plotted. Graph A denotes values for grocery stores, Graph B for general merchandise retailers, Graph C for high-quality clothing and shoes retailers, Graph D for low-quality clothing and shoes retailers, and Graph E for all clothing and shoes retailers, and Graph F for restaurants.



However, as a result of changes in the composition of spending across categories, we find that the net effect on the average labor intensity of the bundle of retail spending we observe is reversed. For instance, as income increases, households shift spending away from clothing retailers and toward general merchandise. Because clothing stores tend to be much more labor intensive than general merchandise stores, the decrease in average labor intensity driven by this cross-category substitution dwarfs the within-category trend toward higher-quality and more

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

labor-intensive retailers. Although there is strong evidence that substitution across retailers can significantly amplify the labor-demand component of the business cycle *within* a range of industries, understanding how household spending shifts *across* all sectors, not solely retail, is necessary to fully characterize the entire effect.

V. Conclusion

Using transaction-level financial data across hundreds of thousands of households, we illustrate several dimensions of heterogeneity in household firm choice in the context of retailers.

First, we document new facts about the distribution of retailer choice, both in the cross section of households and also within households, as income fluctuates. We find that households with higher incomes spread retail spending across a wider range of firms, patronize smaller retailers more often, and are more willing to experiment with new retailers.

Income fluctuations among customers represent a substantial risk to firms. We demonstrate that retailer choice is an important channel of adjustment to fluctuations in household income, much the same as has been studied with home production and searching more extensively for lower prices. Households shift the composition of retailers toward higher-quality and more local retailers as income increases, and households tend to patronize retailers of lower quality as the time since their last paycheck increases. This last trend is especially true for households that exhibit a high MPC over time.

Finally, we show that substitution across retailers is related to firm characteristics and thus can affect both the risk borne by individual firms and also broader business-cycle dynamics. Although these shifts are not all quantitatively large, the fact that these changes do not average out across different retailers implies that retailer choice affects key financial and macroeconomic outcomes, such as profitability and labor demand.

Our study is limited in scope to the set of large customer-facing retailers because they provide the most accessible set of firms interacting with consumers through direct financial transactions. We think that equivalent shifts across suppliers and customers may occur further up the supply chain within the set of wholesalers and business-to-business firms. However, given the constraints imposed by our household-centric data, we leave the formal analysis of these relationships for future research.

Supplementary Material

To view supplementary material for this article, please visit http://dx.doi.org/10.1017/S0022109020000526.

References

Adams, R.; H. Almeida; and D. Ferreira. "Understanding the Relationship between Founder–CEOs and Firm Performance." *Journal of Empirical Finance*, 16 (2009), 136–150.

Ait-Sahalia, Y.; J. A. Parker; and M. Yogo. "Luxury Goods and the Equity Premium." *Journal of Finance*, 59 (2004), 2959–3004.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

Baker, S. R. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy*, 126 (2018), 1504–1557.
Baker, S. R., and C. Yannelis. "Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown." *Review of Economic Dynamics*, 23 (2017), 99–124.
Bils, M., and P. J. Klenow. "Quantifying Quality Growth." *American Economic Review*, 91 (2001), 1006–1030.
Blinder, A.; E. R. Canetti; D. E. Lebow; and J. B. Rudd. *Asking about Prices: A New Approach to Understanding Price Stickiness*, New York, NY: Russell Sage Foundation (1998).
Bustamante, M. C., and A. Donangelo. "Product Market Competition and Industry Returns." *Review of Financial Studies*, 30 (2017), 4216–4266.
Campbell, J. Y.; M. Lettau; B. G. Malkiel; and Y. Xu. "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk." *Journal of Finance*, 56 (2001), 1–43.
Coibion, O.; Y. Gorodnichenko; and G. H. Hong. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review*, 105 (2015), 993–1029.
Davis, S. J.; J. Haltiwanger; R. Jarmin; and J. Miranda. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." *NBER Macroeconomics Annual*, 21 (2006), 107–179.
Davis, S. J., and J. Kahn. "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels." *Journal of Economic Perspectives*, 22 (2008), 155–180.
Deaton, A. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore, MD: Johns Hopkins University Press (1997).
Deaton, A. "Measuring and Understanding Behavior, Welfare, and Poverty." *American Economic Review*, 106 (2016), 1221–1243.
Dupraz, S., *A Kinked Demand Theory of Price-Rigidity*, Ph.D. thesis, Columbia University (2016).
D'Acunto, F.; R. Liu; C. Pflueger; and M. Weber. "Flexible Prices and Leverage." *Journal of Financial Economics*, 129 (2018), 46–68.
Einav, L.; P. J. Klenow; B. Klopack; J. D. Levin; L. Levin; and W. Best. "Assessing the Gains from E-Commerce." Working Paper, National Bureau of Economic Research (2018).
Eisfeldt, A. L., and D. Papanikolaou. "Organization Capital and the Cross-Section of Expected Returns." *Journal of Finance*, 68 (2013), 1365–1406.
Fama, E. F., and K. R. French. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, 47 (1992), 427–465.
Fama, E. F., and K. R. French. "New Lists: Fundamentals and Survival Rates." *Journal of Financial Economics*, 73 (2004), 229–269.
Fama, E. F., and K. R. French. "Dissecting Anomalies." *Journal of Finance*, 63 (2008), 1653–1678.
Foster, L.; J. Haltiwanger; and C. Syverson. "The Slow Growth of New Plants: Learning about Demand?" *Economica*, 83 (2016), 91–129.
Ganong, P., and P. Noel. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review*, 109 (2019), 2383–2424.
Gomes, J. F.; L. Kogan; and M. Yogo. "Durability of Output and Expected Stock Returns." *Journal of Political Economy*, 117 (2009), 941–986.
Gorodnichenko, Y., and M. Weber. "Are Sticky Prices Costly? Evidence from the Stock Market." *American Economic Review*, 106 (2016), 165–199.
Hall, R. E. "General Equilibrium with Customer Relationships: A Dynamic Analysis of Rent-Seeking." Working Paper, Society for Economic Dynamics (2008).
Herskovic, B. "Networks in Production: Asset Pricing Implications." *Journal of Finance*, 73 (2018), 1785–1818.
Herskovic, B.; B. Kelly; H. Lustig; and S. Van Nieuwerburgh. "Firm Volatility in Granular Networks." Working Paper, National Bureau of Economic Research (2017).
Jaimovich, N.; S. Rebelo; and A. Wong. "Trading Down and the Business Cycle." *Journal of Monetary Economics*, 102 (2019), 96–121.
Jappelli, T., and L. Pistaferri. *The Economics of Consumption: Theory and Evidence*, New York, NY: Oxford University Press (2017).
Kueng, L. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics*, 133 (2018), 1693–1751.
Mastrobuoni, G., and M. Weinberg. "Heterogeneity in Intra-Monthly Consumption Patterns, Self-Control, and Savings at Retirement." *American Economic Journal: Economic Policy*, 1 (2009), 163–189.
Nevo, A., and A. Wong. "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession." *International Economic Review*, 60 (2018), 25–51.
Olafsson, A., and M. Pagel. "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software." *Review of Financial Studies*, 31 (2018), 4398–4446.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press

Pan, Y., and G. M. Zinkhan. "Determinants of Retail Patronage: A Meta-Analytical Perspective." *Journal of Retailing*, 82 (2006), 229–243.
Rotemberg, J. J. "Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy." *Journal of Monetary Economics*, 52 (2005), 829–852.
Shapiro, J. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics*, 89 (2005), 303–325.
Stephens, M. "3rd of the Month: Do Social Security Recipients Smooth Consumption between Checks?" *American Economic Review*, 93 (2003), 406–422.
Stroebel, J., and J. Vavra. "House Prices, Local Demand, and Retail." *Journal of Political Economy*, 127 (2019), 1391–1436.
Van Binsbergen, J. H. "Good-Specific Habit Formation and the Cross-Section of Expected Returns." *Journal of Finance*, 71 (2016), 1699–1732.

https://doi.org/10.1017/S0022109020000526 Published online by Cambridge University Press