Journal of Financial Economics 151 (2024) 103758

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Financial returns to household inventory management ☆

Scott R. Baker [a], Stephanie Johnson [b,*], Lorenz Kueng [c]

[a] Northwestern University and NBER, 2169 Campus Drive, 60208, Evanston IL, USA
[b] Rice University, 1900 Rice Boulevard, 77005, Houston TX, USA
[c] Swiss Finance Institute, USI Lugano, and CEPR, Via Giuseppe Buffi 13, 6900, Lugano TI, Switzerland

ARTICLE INFO

Dataset link: https://doi.org/10.17632/66w4y6sb6d.1

JEL classification:
G51
G11
D14
D13
D11
E21

Keywords:
Household working capital
Financial returns
Inventory
Stockpiling

ABSTRACT

Households tend to hold substantial amounts of non-financial assets in the form of consumer goods inventories that are unobserved by traditional measures of wealth, about $725 on average for products covered by our sample. Such holdings can eclipse total financial assets among households in the lowest income quintile. Households can obtain significant financial returns from strategically shopping and managing these inventories. In addition, they choose to maintain liquid savings—household working capital—not just for precautionary motives but also to support this inventory management. We demonstrate that households earn high marginal returns from investing in household working capital, well above 20% at low levels of inventory, though these marginal returns decline rapidly as inventory increases. Nevertheless, average returns from inventory management are high—about 50% for the typical household—and affect household portfolio returns substantially for all but the top income and asset quintiles. We provide evidence from scanner and survey data that supports this conclusion. For many households, working capital is therefore an important asset class that has been largely ignored by the household finance literature, and inventory management provides them with an alternative to investing in risky financial markets at low levels of liquid wealth.

1. Introduction

While a large number of American households hold small amounts or even zero financial assets, all households hold at least some resources in the form of consumer goods inventories. These inventories can be managed over time through strategic shopping behavior as households are able to take advantage of coupons, temporary low prices at retailers, and savings from buying in bulk. In this paper, we study how the financial return to investment in inventories affects households' desire to hold liquid assets like cash and cash equivalent assets (e.g., checking accounts, transaction accounts, credit card lines of credit). We refer to these combined resources—the sum of cash and inventory—as *household working capital*. This combination of financial resources and consumer goods echoes firms' working capital which includes both current account resources as well as materials and inventories that may be at least partially non-tradeable.

The paper makes two main contributions. First, using scanner data from NielsenIQ and income and asset data from the Survey of Consumer Finances (SCF), we quantify this hitherto neglected source of non-financial wealth on households' balance sheets using a new method to impute inventory from flows of NielsenIQ goods. Aggregating across all NielsenIQ goods included in our sample, we find that households

☆ Toni Whited was the editor for this article. We thank the editor, two anonymous referees, Alex Butler, Hector Calvo-Pardo, James Choi, Giuseppe Di Giacomo, John Gathergood, Gustavo Grullon, Luigi Guiso, Michael Haliassos, Xavier Jaravel, Pierre Mabille, Soren Leth-Petersen, Jonathan Parker, Luigi Pistaferri, Tarun Ramadorai, Antoinette Schoar, and Qingyuan Yang for valuable feedback and our discussants Francesco d'Acunto, Andres Almazan, Matteo Benetton and Pascal Noel and participants at seminars at Einaudi Institute for Economics and Finance (EIEF), National University of Singapore (NUS), Northwestern University, Rice University, CEPR European Conference on Household Finance 2020, Copenhagen New Consumption Data Conference, Western Finance Association Meeting 2021, and NBER Household Finance Summer Institute 2022 for their comments. Researchers own analyses calculated (or derived) based in part on data from The Nielsen Company (US), 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 researchers 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.

* Corresponding author.
E-mail address: stephanie.g.johnson@rice.edu (S. Johnson).

https://doi.org/10.1016/j.jfineco.2023.103758
Received 22 November 2021; Received in revised form 7 November 2023; Accepted 7 November 2023
Available online 24 November 2023

hold on average around $725 in consumer goods inventory at any given time. For an average household in the lowest quintile of annual household income (under $22,000), inventory represents a greater store of value than their total household financial assets.

Second, we build a parsimonious model of inventory management to compute the marginal financial return to investing in household working capital through the maintenance of liquid savings and engaging in strategic shopping behavior. These marginal returns, net of product depreciation and trip costs, are household specific, scale with consumption, are approximately risk-free, and are above 20% at low levels of working capital.

The model highlights two key sources of returns. By taking larger and less frequent trips, households can save on trip fixed costs and also take advantage of lower unit prices by buying goods in bulk. Alternatively, consumers can shop more frequently, giving them additional opportunities to take advantage of temporary deals at retailers but at higher cumulative trip fixed costs. This stockpiling of non-durable but storable products can produce volatile expenditures alongside smooth consumption paths, mirroring a tendency recognized among durable goods (Parker 1999; Browning and Crossley 2001, 2009).

Both strategies require a substantial amount of resources: liquid assets in the former to pay for the larger trip sizes, and consumer inventory in the latter, which is associated with depreciation costs. The household in the model optimally chooses shopping trip frequency to minimize the cost of providing a given consumption stream, subject to a household working capital constraint. The model therefore allows us to study how investing in household working capital generates a return in the form of reduced trip costs and lower per unit prices, taking into account product depreciation.

Beyond large marginal returns, average returns from inventory management are also large for all households, well above 30% even at high levels of working capital, although both marginal and average returns decline with income and financial assets. Including working capital in the household balance sheet increases household portfolio returns substantially both because of the high average returns and because working capital has a high portfolio share for many households. Moreover, because lower-income households hold a larger share of wealth in the form of household inventory, including working capital more than offsets the higher returns richer households achieve on their financial assets. Total returns, including working capital, decline as households' income and financial assets increase.

For a minority of consumers, the high marginal returns observed in our data can even rationalize high cost borrowing to finance inventory purchases (e.g., credit cards). Households in our sample generally engage in substantial amounts of inventory management and thereby likely have low *marginal* returns for additional working capital accumulation. Even for these households that have exhausted the potential for profitable additional investment, *average* returns from optimal inventory management remain high (around 50% at the median ratio of inventory to annual spending). Optimal inventory management therefore provides a rationale for why households hold sizable amounts of household working capital above and beyond the desire to maintain a buffer stock or precautionary source of savings. For instance, Orhun and Palazzolo (2019) note that low-income households are less responsive than higher income households to sales or promotions in part due to a lack of liquidity reserves to employ for intertemporal substitution.

Existing models of deal shopping focus on individual products in a stochastic framework (Boizot et al. 2001; Hendel and Nevo 2013). In contrast, we focus on the deterministic steady state that results from aggregating all NielsenIQ products a household purchases, where a constant fraction of goods is on sale at any given time across a household's total basket. This formulation is derived from an assumption of independent price deals across goods and backed by observations from the data. It has implications for households' cash holdings. In particular, if deals are approximately independent across products at typical shopping trip frequency, stocking up in response to deals is consistent with a deterministic steady state where consumers hold a substantial level of inventory at all times, but where trips are consistently spaced and of a similar size.[1]

Our model builds on a previous literature which shows that consumers use stockpiling strategically to take advantage of temporarily low prices and to reduce the frequency of shopping trips. While our model is static, much of that previous literature has exploited temporary shocks to causally identify the stockpiling channel of consumer responses to these shocks. Baker et al. (2019) and Baker et al. (2021) use anticipated local sales tax increases at a monthly frequency and show that consumers respond strongly along several margins, including stocking up on products while taxes are still low. Hendel and Nevo (2006a) identifies the related effect of temporary price discounts on inventories in the context of a dynamic discrete choice model of individual product demand with exogenously given shopping trip frequency. The study underscores the importance of household inventory to explain the large price elasticities observed in scanner data.

By computing the total net returns to investment in household working capital, we go one step further than existing work, which focuses on in-store savings as a percentage of the product price, but does not take into account the additional household working capital that must be held to facilitate these savings and the financial returns to this working capital (Griffith et al. 2009; Nevo and Wong 2019). We also extend our framework to include the costs from product depreciation and spoilage and the relation between the level of inventory holdings and differences in shopping trip fixed costs associated with different shopping behaviors. Moreover, relative to the previous literature, we endogenize the timing of shopping trips.

By highlighting the role of household working capital for households' portfolio allocation and spending behavior—especially for households with relatively low income and low financial wealth—our paper relates to a large literature in household finance. While inventories have long been recognized as an important part of *firms'* working capital and have received considerable attention in finance (Petersen and Rajan 1997; Fisman and Love 2003; Yang and Birge 2018; Rampini 2019), inventories of consumer goods and household working capital have been largely ignored by the household finance literature.[2]

For instance, none of the country studies of household portfolios in the widely cited book by Jappelli et al. (2002) include household inventories. This also applies to the chapter by Bertaut and Starr (2000), which studies U.S. households' portfolios. One explanation for this gap is that inventories are often difficult to observe and measure. For example, they are missing from traditional consumer finance data such as the SCF. In addition to quantifying gross and net financial returns to household inventory management, our second main contribution is to quantify the level of working capital and its distribution across households. Our paper is therefore one of the first systematic studies of the role of household inventories in household finance.[3]

[1] In contrast, if aggregated deals are autocorrelated, households may want to hold substantial additional cash to stock up more in those (random) weeks. We provide empirical evidence supporting the relative independence of deals across products over time at typical shopping trip frequency. With few exceptions (e.g., Black Friday, New Year's or Boxing Day sales), retailers generally feature consistent amounts of goods on sale throughout the year rather than concentrating deals in particular weeks. Appendix Figure H.1a ranks retailer-weeks by deal share relative to the retailer average. For a given retailer, the weekly deal share varies from around 80% to 120% of the annual average, with most weeks having fairly similar deal shares. To the extent that there are fluctuations in deal share, these do not seem to be strongly correlated across retailers. Appendix Figure H.1b shows that, pooling across retailers, there is no calendar week with particularly high or low deal share relative to the mean.

[2] A notable exception is Samphantharak and Townsend (2010), which focuses on households in developing economies who are engaged in agriculture and therefore have a large fraction of their wealth invested in inventories.

[3] There is a literature in macroeconomics studying heterogeneity in the effective price paid for similar goods across households and over business cycles

Finally, we highlight some of the differences between household inventory and firm inventory as they are not identical. One key difference between household and firm inventory is that firm inventory is held for sale. While households may occasionally re-sell items, this is not typically the intended purpose of household inventory. In traditional firm inventory models, firms minimize the cost of meeting customer demand. In a model of household inventory, households manage inventory to minimize the cost of supplying their personal consumption.

The remainder of the paper is structured as follows. Section 2 describes the data sources. Section 3 discusses how we construct our measure of household inventory. Section 4 lays out the household shopping model. Section 5 computes the financial net return to investing in household working capital and tests some predictions of the model. Section 6 concludes.

2. Data

Our analysis uses data from five main sources, the NielsenIQ Consumer Panel (NCP), the NielsenIQ Retail Scanner Panel (NRP), the Survey of Consumer Finances (SCF), the Food Safety and Inspection Service Foodkeeper Data (FSIS), and the National Health and Nutrition Examination Survey (NHANES).

2.1. NielsenIQ consumer panel (NCP)

The Nielsen Company Consumer Panel (2013–2014) consists of a long-run panel of over 60,000 nationally representative American households in 52 metropolitan areas. Using bar-code scanners and hand-coded diary entries, participants are asked to report all spending on household goods that they engage in and also to detail information about the retail location that they visited in a given trip. NielsenIQ uses monetary prizes and continual engagement with panelists to try to maintain high levels of continued participation and limit attrition ($\leq$ 20% per year) from the sample. On average, we observe $306 of spending per month for each household on covered product groups.[4]

The NCP is constructed to be a representative sample of the US population. Broda and Weinstein (2010) and Einav et al. (2010) perform more analysis of the NCP. Overall, they deem the NCP to be of comparable quality to many other commonly-used self-reported consumer data. The NCP primarily covers trips to grocery, pharmacy, and mass merchandise stores but also spans a wider range of channels such as catalog and online purchases, liquor stores, delis, and video stores.

In this paper, we utilize data from the 2013 and 2014 NCP unless otherwise noted. Our measure of household inventory is necessarily limited by the scope of the NCP. To the extent that households stockpile clothing, electronics or other larger purchases, we will underestimate inventory and thus consider our values a lower bound of household inventory.

2.2. NielsenIQ retail scanner panel (NRP)

The Nielsen Company Retail MSR Scanner Data (2013–2014) contains price and quantity information at the store-week level of each UPC carried by a covered retailer. This data covers almost 100 retail chains with over 40,000 unique stores in over 350 MSAs across the country.

In general, the data span many of the largest retailers in the grocery, mass merchandiser, drugstore, and pharmacy sectors. Within the store, the data provide a comprehensive view of products sold, with more than two million unique product identifiers (i.e., scanner codes or UPCs) across 1,305 product modules, 118 product groups, and 10 departments. During these years, the database picks up about half of total sales in grocery stores and pharmacies and about 30% of sales in other retailers.

2.3. Survey of consumer finances (SCF)

The Survey of Consumer Finances (2010, 2013, 2016) of the Board of Governors of the Federal Reserve System contains detailed information on U.S. households' income and assets. Income is gross household income over the calendar year preceding the survey. Financial assets include checking accounts, savings accounts, CDs, mutual funds, bonds, stocks, and money market funds.

2.4. USDA food safety and inspection service foodkeeper data (FSIS)

The Food Safety and Inspection Service FoodKeeper Data (2020) of the U.S. Department of Agriculture contains information on recommended food and beverage storage times. We rely primarily on this information to infer depreciation estimates for each NielsenIQ product module.

2.5. CDC national health and nutrition examination survey (NHANES)

To provide direct empirical evidence on households' actual consumption of the products we consider, we look at the CDC's National Health and Nutrition Examination Survey (2013–2014). Survey respondents report food and beverage items consumed on two non-consecutive days, with the second day being 3–10 days after the first day. We restrict attention to items households purchased in grocery stores, supermarkets, or convenience stores, as this most closely corresponds to purchases covered by NielsenIQ.

We manually assign each of over 4,000 food items to a NielsenIQ product group. The item information is very detailed, for example distinguishing between whether a vegetable item was canned, fresh or frozen. In some cases the NHANES code corresponds to a meal with multiple ingredients (e.g., "Frankfurter or hot dog sandwich, beef, plain, on wheat bun"). In this case, we assign multiple product group codes. The NHANES codes broadly correspond to a level of aggregation between UPC and NielsenIQ 'Product Module'.

3. Measuring levels of household inventory

Although we can track the flow of purchases for individual products over time in the NCP, including the exact time stamp of each purchase, we must make three assumptions to compute household inventories because we do not observe initial inventory or the flow of consumption.

Our first assumption is that a household has just enough in stock initially to ensure that inventory does not become negative at any point during the year. This implies that inventory hits zero at least once each year and therefore underestimates inventory if households violate this assumption. The second assumption is that annual inventory depletion equals annual spending. The third assumption is that a given product's

(Chevalier et al. 2003; Aguiar and Hurst 2007; Coibion et al. 2015; Kaplan and Menzio 2016; Kaplan and Schulhofer-Wohl 2017; Stroebel and Vavra 2019).

[4] We exclude throughout the analysis product modules for which we believe that either NielsenIQ would not provide good coverage, or that our assumptions are unlikely to hold. We exclude all product modules within the following product groups: "Tobacco", "Ungrouped Items", "Hardware", "Housewares", "Toys and Sporting Goods", "Seasonal", "Beer", "Wine", "Liquor". We also exclude the following product modules: "Cellular phone", "Computer Software", "Printers", "Video Products Prerecorded", "Video and Computer Games", "Computer Software and Supply", "Telephone and Accessory", "Camera", "Paper Shredders", "Prepaid Gift Cards". The vast majority of spending in "Ungrouped Items" is accounted for by gas and apparel. Excluded products account for 14.1% of NielsenIQ spending on average, and 9.4% of spending for the median household.

rate of inventory depletion (consumption and depreciation) is constant throughout the year. Using direct evidence from NHANES on actual consumption rather than spending, Appendix B shows that consumption is indeed fairly constant when aggregating up to NielsenIQ product groups, which is our preferred level of aggregation (see also Aguiar and Hurst (2013) for similar evidence). With these three assumptions, we can compute the initial level of inventory for each household in the NCP sample as well as the inventory level at all later points in time.

How well these assumptions recover unobserved true household inventories depends on the level of product aggregation. To build intuition, it is useful to think of two extreme cases. On the one hand, if we do not aggregate individual products (UPCs) at all, then our assumption that consumption is constant throughout the year is poor and inventory is overstated. For instance, if a household switches cereal products sequentially to try more varieties, say weekly between Kellogg's Raisin Bran Original and Kellogg's Raisin Bran Crunch, they do not hold a stockpile of all varieties at once, but rather consume them one after the other. If the household consumes the same amount of cereal every day, consumption is constant at the product *group* level (cereal) but not at the product level. Assuming constant consumption of each product would overstate inventory because we would incorrectly infer that the household consumed stockpiled Raisin Bran Original in weeks where we do not observe such a purchase, whereas in reality they just consumed Raisin Bran Crunch in that week.

On the other hand, choosing a very broad level of aggregation can lead to inventory being understated. The broader the product categories used, the less likely it is that the household completely runs out of each product category at some point during the year. For example, a household may run out of canned tomatoes at some point during each year, but may only rarely have a pantry completely empty of all canned goods. If households' true inventories at the assumed level of aggregation do not hit zero at some point during the year, our imputed initial level inventory will be too low.

For these reasons, the average level of measured inventory varies with the level of product aggregation we assume. To show the effect of aggregation, we impose our three assumptions for each individual product category as well as for individual products (i.e., no aggregation) and sum up inventories over all categories to get total household inventory. Appendix Table H.1 provides summary statistics of household inventory by aggregation.

With this approach, the average amount of household inventory in NCP goods aggregated to our preferred level of 118 NielsenIQ "product groups" (e.g., "cheese") is $725. Aggregating to the 1,305 NielsenIQ "product modules" (e.g., "natural American cheddar") yields a slightly higher household inventory of $985 while aggregating to a UPC level gives an average inventory of $1,461. In our opinion, this likely overstates inventories substantially as the constant consumption assumption is inappropriate for individual products. On the other hand, aggregating aggressively to the 10 NielsenIQ "departments" (e.g., "dairy") produces an average inventory of $431, which very likely underestimates true inventory because it likely never reaches zero during the year for most departments. In the remainder of the paper we therefore aggregate UPCs to "product groups", and we document in Appendix H how our main findings change for different levels of product aggregation.

Next, we derive the formula for average inventory, explain how we implement this formula with the data at hand, and provide two validation exercises for our new measure of household inventory.

3.1. Computing inventories

We derive a formula for average household inventory using a value-based approach, with dollars of spending measuring the inflow of inventory.[5,6] The average inventory held over the period from time zero to T is $\bar{I}_T = \frac{1}{T}\int_0^T I(t)dt$, where $I(t)$ is the unobserved level of inventory at time t. Inventory at time t reflects the initial time zero level of inventory $I(0)$, purchases made on trips between time 0 and time t, and the rate of inventory depletion d, which we assume to be constant[7]:

$$I(t) = I(0) + \sum_{j=1}^{n_t} X_{t_j} - d \cdot t, \tag{1}$$

where $\{t_j\}_{j=1}^{n_t}$ are the dates of the n_t shopping trips occurring between time 0 and time t, corresponding to the time stamps in the NCP data. X_{t_j} is total expenditures on the jth trip. $I(0)$ contributes to average inventory from time 0 to T and each trip size X_{t_j} contributes from time t_j to T so that the integral is:

$$\int_0^T I(t)dt = I(0) \cdot T + \sum_{j=1}^{n_T} (T - t_j) X_{t_j} - \frac{T^2}{2} d. \tag{2}$$

Hence, average inventory is:

$$\bar{I}_T = I(0) + \sum_{j=1}^{n_T} \frac{T - t_j}{T} X_{t_j} - \frac{T}{2} d. \tag{3}$$

We compute average annual inventory in the data by applying (3) to each product group g for each household h. Hence, with time measured in years, we have $T=1$ and the time stamp t_j of trip j is relative to the start of the year and takes values between 0 and 1 (e.g., $t_j=0.5$ corresponds to June 30). Assuming annual depletion (d) is equal to annual spending during calendar year y ($\sum_{j=1}^{n_y} X_{t_j}$), and adding subscripts h, g, and y to indicate the sources of heterogeneity to which we apply equation (3), annual average inventory of household h in product group g in calendar year y is:

$$\bar{I}_{y,h,g} = I(0)_{y,h,g} + \sum_{j=1}^{n_{y,h}} (1 - t_j) X_{t_j,y,h,g} - \frac{1}{2} \sum_{j=1}^{n_{y,h}} X_{t_j,y,h,g}, \tag{4}$$

where $n_{y,h}$ is the number of trips household h makes over calendar year y (replacing n_T in equation (3)), $X_{t_j,y,h,g}$ is the value of purchases made on trip j in product group g (generalizing X_{t_j}), and $1 - t_j$ is the share of the calendar year remaining when trip j occurs.

[5] Under this approach, fluctuations in unit prices over time may in principle affect the accuracy of our inventory calculation. An alternative approach is to derive average inventory in terms of quantity, and then apply a household-specific average annual per unit price to value it. Given that both approaches yield very similar results on a consistent set of UPCs as shown in Appendix E, we prefer the value approach as it allows use of all NCP products, including those where the unit of measurement is not comparable across items (e.g., products that are not measured in ounces, such as boxes of tissues).

[6] There are several established methods for valuing firm inventory, such as First-In First-Out (FIFO), Last-In First-Out (LIFO), and Weighted Average Cost (WAC). Our approach to valuing household inventory is conceptually closest to WAC, though it is not identical. We value inflows at purchase price and outflows at the weighted average cost of items in the same product group purchased by the household in the same calendar year. We also obtain similar results valuing both inflows and outflows at weighted average cost. Finally, we note that although not standard for business inventory of physical goods, financial assets are marked to market (when classified as available-for-sale). When valuing household inventory, it is important to keep in mind that households do not usually hold inventory with the intention of selling it, and so mark-to-market valuation is unlikely to be appropriate in this case.

[7] Depletion rate d broadly corresponds to the sum of consumption and depreciation in the model of Section 4. The main difference is that in the model we make various assumptions about the depreciation profile depending on the type of item. For some goods, exponential depreciation is a better approximation (see Appendix C). Here, we assume inventory is depleted linearly for simplicity.

We recover the unobserved $I(0)_{y,h,g}$ as the level of initial inventory needed to ensure that inventory of household h in product group g is never negative at any point in year y. To do this, we first compute the inventory remaining immediately prior to each trip j, assuming constant depletion and initial inventory equal to zero. We then find the minimum value of inventory and set:

$$I(0)_{y,h,g} = -\min_j I(t_j)_{y,h,g}. \tag{5}$$

$I(t_j)_{y,h,g}$ is the value of inventory remaining immediately prior to trip j, assuming that $I(0)_{y,h,g} = 0$. That is:

$$I(t_j)_{y,h,g} = \begin{cases} -t_j \cdot \sum_{j=1}^{n_{y,h}} X_{t_j,y,h,g} & \text{if } j = 1, \\ I(t_{j-1})_{y,h,g} + X_{t_{j-1},y,h,g} \\ -(t_j - t_{j-1}) \sum_{j=1}^{n_{y,h}} X_{t_j,y,h,g} & \text{if } j > 1. \end{cases} \tag{6}$$

Note that $I(0)_{y,h,g} \geq 0$, because $I(t_1)_{y,h,g} \leq 0$. We then compute the average aggregate inventory level for household h across all years in the sample (2013–2014) as:

$$\bar{I}_h = \frac{1}{2} \sum_{y=2013}^{2014} \sum_g \bar{I}_{y,h,g}. \tag{7}$$

Fig. 1 shows the distribution of our inventory measure $\bar{I}_h$ across households in terms of both dollars (Fig. 1a) and the ‘inventory ratio’: the value of inventory relative to annual spending on NielsenIQ products (Fig. 1b). The median household holds inventory worth 0.2 years or 2.4 months of spending.

Note that our measure of inventory naturally excludes inventory holdings in goods not covered by the NCP; most notably it excludes all large durable items like cars and furniture. Although the NCP contains some smaller durable items (such as clothing and cookware), we choose to exclude these as we believe our assumptions may be less appropriate in these cases. On the household balance sheet, some of these durable items would be classified as long-term physical assets—corresponding to “Property, Plant, and Equipment (PP&E)” on the corporate balance sheet—and are therefore not included in our definition of household working capital.

Our measure of inventory is also not inflated by product waste. When computing inventory by comparing the timing of spending with the timing of consumption, there would be a concern that the difference reflected not just product storage, but also product depreciation. In practice, we do not observe either consumption or the disposal of spoiled products. Instead, we assume that annual inventory depletion d is equal to annual spending, and we do not need to take a stand on how much of that depletion is consumption and how much is depreciation. In the model of Section 4, we will use FSIS data to calibrate the product depreciation rates in our model, which allows us to explicitly take spoilage into account when computing financial returns.

3.2. How important is inventory in households’ portfolios?

Overall, inventory is an important asset for many households, even with durables excluded. To show this, in Fig. 2 we impute inventory values for each SCF household using characteristics observable in both NielsenIQ and the SCF such as income, age, home prices, and household size and composition.

Then, using data from the 2010, 2013 and 2016 SCF, we compute the inventory portfolio share for each household i, $\frac{\text{Inventory}_i}{\text{Financial Assets}_i + \text{Inventory}_i}$. Our measure of financial assets includes checking accounts, savings accounts, CDs, money market accounts, bonds and stocks (both directly held and in mutual funds). Fig. 2a shows the average and median inventory portfolio share by income quintile and Fig. 2b shows the average value of inventory in each income quintile in the NielsenIQ data. For households in the bottom income quintile

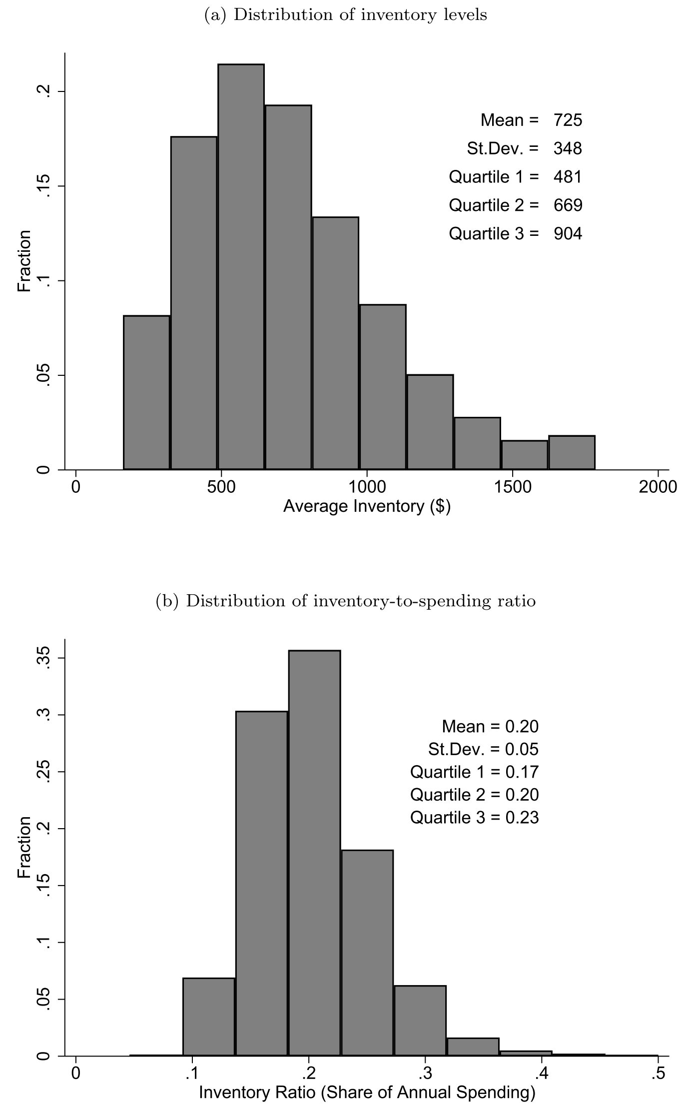

Fig. 1. Observed consumer goods inventory.
Notes: Panel (a) plots how the average 2013-2014 inventory level $\bar{I}_h$ varies across households in the NCP. Average inventory is plotted up to the 99th percentile. Summary statistics reported in the top right corner are computed using all observations. Panel (b) plots the distribution of inventory ratio, i.e., inventory as a share of annual household spending on goods covered by NielsenIQ. Both panels are constructed using NielsenIQ sampling weights. Appendix Table H.1 provides corresponding summary statistics for alternative product aggregation levels (UPC, Product Module, Product Group, and Department).

(up to $22,000), inventories constitute about 60% of the portfolio. As income increases, inventory holdings grow more slowly than financial assets and the inventory portfolio share declines.

3.3. Validating the household inventory measure

We perform two primary validation exercises for our measure. First, we show that our measure of inventory correlates positively with measures of product life. Second, we show that households run down inventories in advance of a move. While this evidence is not conclusive, it shows that our measure of inventory is associated with the properties and household behavior we would expect.

3.3.1. Measured inventory and product life

We compute inventory for each household-product group combination and estimate the relationship between product group durability

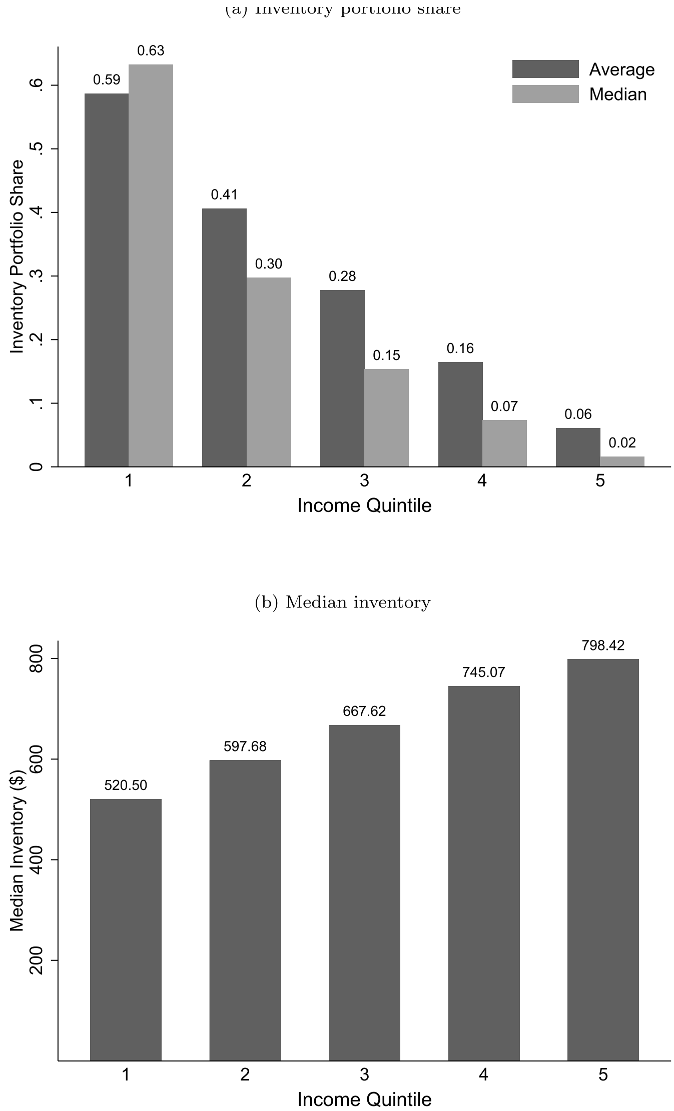

Fig. 2. Inventory portfolio share by income.

Notes: Panel (a) is constructed by combining data from the NCP over 2013 and 2014 and the SCF over 2010, 2013, and 2016. We impute inventory for SCF households based on characteristics observable in both datasets: house price, household income, the maximum age of household members, household size, marital status, and indicators equal to one if the respondent identifies as non-Hispanic white, if the household contains young children, if all adults work full time, or if either respondent or spouse has a college degree. We train a machine learning model to predict inventory using the Matlab command *fitrensemble* with hyperparameter optimization. The resulting method is LSBoost with 11 trees, a learn rate of 0.35 and a minimum leaf size of 46. The model explains 17.2 per cent of variation in observed inventory out-of-sample. Then, using data from the 2010, 2013 and 2016 SCF, we compute the inventory portfolio share for each household i, $\text{Inventory}_i/(\text{Financial Assets}_i + \text{Inventory}_i)$, and report the average and median share by income quintile. Financial assets include checking accounts, savings accounts, CDs, money market accounts, bonds and stocks (both directly held and in mutual funds). We do not subtract debt and do not include retirement accounts (see Appendix Figure H.2 for results including retirement accounts). Income is reported to the nearest thousand dollars. Panel (b) shows the median value of inventory in each income quintile computed using the NCP. The lower cutoffs for each income quintile are $0, $22,000, $38,000, $61,000, and $101,000. We use NielsenIQ sampling weights in both panels.

and inventory. Table 1 shows that inventory as a share of NielsenIQ spending is increasing in durability. It serves as a check on the magnitudes for our calculations of inventory levels. To measure durability, we manually assign each NielsenIQ product module a usable life in months (between 0.03 and 60) using FSIS data. The majority of spending in our NCP sample is on products with a lifetime of more than 6 months.

The relationship between inventory and shelf life is increasing and concave. It is robust to controlling for the household's trip frequency, which also has substantial explanatory power for the inventory ratio. It is also robust to using within household variation in shelf-life across product groups. Starting from a shelf life of zero, a one month increase in shelf life raises the inventory ratio by around 0.4 percentage points. Column 5 shows an alternative specification with indicators for product groups with an average shelf life greater than 6 months and less than 0.58 months. A shelf life < 0.58 months corresponds to the two most perishable products in the model in Section 4 ($l=1$ and $l=4$). Shelf life ≥ 6 months corresponds to the most storable product ($l=4$). Households hold around 5 percentage points of annual spending more inventory of product groups with an average shelf-life of at least 6 months, and around 2.5 percentage points less of product groups with an average shelf-life less than 0.58 months.

3.3.2. *Inventory dynamics of movers*

Because it is costly to transport a large stockpile of consumer goods, we expect that households will adjust their stockpiling behavior around the time they move. Specifically, we anticipate that households will run down their stockpile prior to a move and therefore reduce purchases.

Fig. 3 shows that the behavior of the subset of households who move is consistent with this conjecture. Households cut spending substantially well in advance of a move, consistent with our finding that they hold a large stockpile of inventories. Spending returns to normal immediately following the move. An alternative explanation for the decline in purchases prior to a move is that households are cutting consumption to pay for move-related expenses—for example to cover transportation costs, down payments, or security deposits. However, we also find that the share of coupon or deal purchases declines around the move (see Appendix Figure H.11), which seems inconsistent with this interpretation. We discuss this further in Section 5.

Assuming that the decline in spending reflects households running down inventory, we can interpret the cumulative decline in spending as a lower bound for steady-state inventory. We expect that some households may transport part of their stockpile—for example if an employer is paying for the move, or if they are only moving a short distance. Hence, the observed decline is likely much smaller than total inventory. For households moving to a new 3-digit ZIP Code, the cumulative spending decline is 6.6% of annual spending.[8] Restricting the sample to households moving more than 974 kilometers (the top quartile of move distance) the cumulative decline is 11.5%. This provides additional evidence that households hold at least several hundred dollars of inventory, independent of our assumptions in Section 3. Comparing these results with our headline inventory valuation, the cumulative decline for long distance movers suggests that the inventory ratio is *at least* 12% on average (compared with 20% computed in Section 3.1).

In Appendix Figure H.3 we repeat the exercise using log quantity purchased instead of log spending. While the use of quantities limits us to items measured in ounces, it allows us to address the fact that spending will tend to understate the pre-move decline in inventory because households take less advantage of deals during this period and therefore pay higher prices. We indeed find that quantities decline slightly more than spending, 7% across all movers and 12.2% for long distance movers.

3.4. *Heterogeneity in inventory holdings across households*

In the previous sections, we demonstrate that consumer goods inventory varies substantially across households. In Table 2, we examine

[8] Using the imputation estimator proposed by Borusyak et al. (2021) gives very similar results; see Appendix Figure H.4.

Table 1
Validation: relationship between durability and inventory ratio.

	(1)	(2)	(3)	(4)	(5)
Shelf Life (Months)	0.128***	0.395***	0.401***	0.392***	
	(0.001)	(0.004)	(0.003)	(0.004)	
Shelf Life Squared		-0.004***	-0.004***	-0.004***	
		(0.000)	(0.000)	(0.000)	
Avg. # Days Between Trips				0.467***	
				(0.008)	
Group Shelf Life ≤ .58 months					-2.220***
					(0.027)
Group Shelf Life > 6 months					5.234***
					(0.027)
Household FE			X		X
Number of Observations	5,578,528	5,578,528	5,578,527	5,535,390	5,578,528
Adjusted R-squared	0.07	0.09	0.27	0.13	0.25

Notes: This table combines data from the NCP over 2013 and 2014 and the FSIS. We estimate variations on the following regression specification, where h indexes households and g indexes NielsenIQ product groups:

$$\text{Inventory Ratio}_{h,g} = b_0 + b_1 \text{Shelf Life}_g + b_2' X_h + e_{h,g}.$$

$\text{Inventory Ratio}_{h,g}$ is the ratio of household inventory to annual spending in product group g, multiplied by 100. Columns 3 and 5 include household fixed effects. Standard errors are clustered by household. Regressions are weighted, using NielsenIQ sampling weights multiplied by total product group expenditures. * $p < .1$, ** $p < .05$, *** $p < .01$.

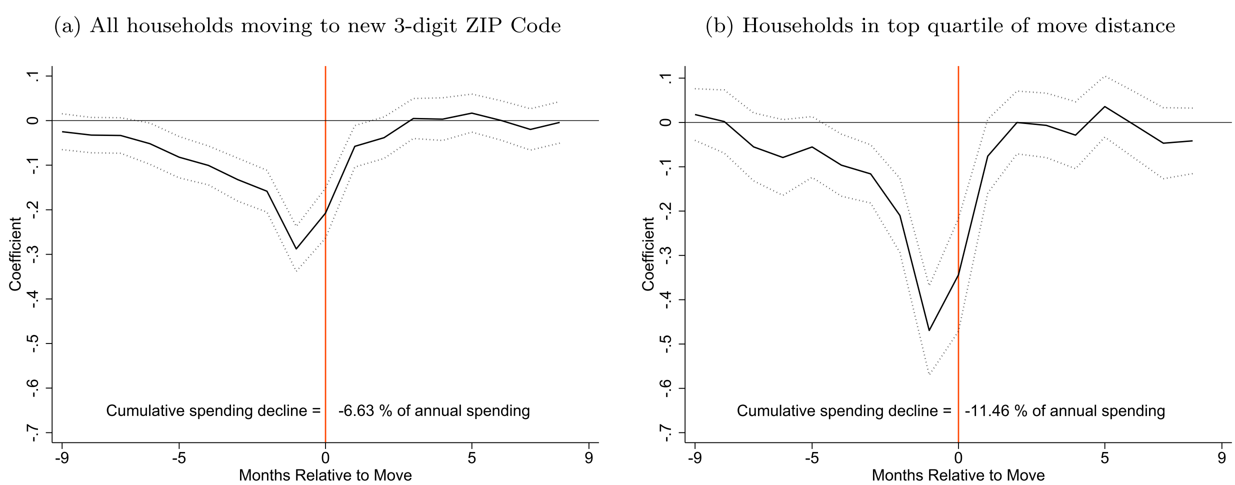

Fig. 3. Validation: spending around move dates.
Notes: This figure shows the change in log spending around the time a household moves. For households who move to a new 3-digit ZIP Code in a given year we impute the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather than their old 3-digit ZIP Code). The figures plot estimates of b_s from the following specification and a 95 per cent confidence interval: $\ln \text{Spending}_{i,t} = \sum_{s=-9}^{9} b_s \text{ Moved}_{i,t-s} + \text{Month FE} + \text{Household FE} + e_{i,t}$, where $\ln \text{Spending}_{i,t}$ is the log spending of household i in month t. $\text{Moved}_{i,t}$ is an indicator equal to 1 if household i moved in month t. The sample includes non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. Panel (b) includes only households in the top quartile of move distance—that is households moving more than 974 km. Standard errors are clustered by household. The regression is weighted using NielsenIQ sampling weights. Appendix Figure H.4 shows robustness to using the imputation approach described by Borusyak et al. (2021) to deal with potential bias in pooled event studies with staggered events.

the extent to which observable characteristics of NielsenIQ households are correlated with their inventory holdings. In Column 1, we note that older households hold less inventory relative to spending, as do married couples and households where all adults work full-time. We also see variation across race and ethnicity, with white households holding less inventory and Asian households holding more. Single person households hold more inventory relative to spending than larger households. Together, these characteristics explain 9 per cent of the variation across households. Column 2 adds financial and housing characteristics. Living in a single family home, having higher income, and living in an expensive or high density ZIP Code all increase inventory relative to spending. Adding these variables has very little effect on explanatory power, but reduces the magnitude of the coefficients on labor force and education variables.

Columns 3 and 4 add characteristics relating to product and store choice. While these characteristics are obviously outcomes of the household's shopping and consumption decision, they have considerable explanatory power and help to form a picture of households who hold high inventory. We also expect these characteristics to be related to fundamental household preferences, which we do not observe directly.

Using data on product expiry dates, we assign each NielsenIQ product module a time to expiration. We classify spending as 'perishable' if the item lasts under about two weeks. In particular, we use a cutoff of 0.58 months which corresponds to the two most perishable goods ($l = 1$ and $l = 2$) in the model. Unsurprisingly, Column 3 shows that households with a high share of perishable spending hold substantially less inventory and perishability explains an additional 8 per cent of variation in the inventory ratio.

Table 2
Factors correlated with household inventory to spending ratios.

	(1)		(2)		(3)		(4)	
Maximum Age (Years)	-0.290	(0.014)	-0.307	(0.015)	-0.296	(0.014)	-0.317	(0.013)
Maximum Age Squared	0.002	(0.000)	0.002	(0.000)	0.002	(0.000)	0.002	(0.000)
Young Children	0.020	(0.111)	-0.005	(0.111)	0.142	(0.106)	0.210	(0.103)
Married	-0.392	(0.074)	-0.596	(0.078)	-0.607	(0.074)	-0.518	(0.072)
All Adults Work Full-time	0.397	(0.066)	0.193	(0.069)	0.213	(0.066)	0.278	(0.064)
White	-1.278	(0.074)	-1.276	(0.074)	-1.205	(0.070)	-0.889	(0.068)
Asian	1.496	(0.182)	1.352	(0.183)	1.555	(0.170)	1.258	(0.164)
Single Household	1.737	(0.085)	1.902	(0.086)	1.913	(0.082)	1.865	(0.081)
College Degree	0.336	(0.059)	0.107	(0.063)	0.264	(0.060)	0.236	(0.058)
Single Family Home			0.594	(0.072)	0.542	(0.069)	0.465	(0.068)
ZIP Code House Price ($00,000s)			0.054	(0.021)	0.114	(0.019)	0.066	(0.018)
Income ($000s)			0.007	(0.001)	0.008	(0.001)	0.005	(0.001)
ZIP Code Persons per Sq. Mi. (000s)			0.013	(0.005)	0.011	(0.004)	0.004	(0.004)
Perishable Share of Spending					-14.647	(0.333)	-15.517	(0.379)
Discount Store Share							0.566	(0.126)
Dollar Store Share							4.697	(0.528)
Drug Store Share							3.794	(0.379)
Convenience Share							2.849	(1.258)
Online Share							4.548	(0.537)
Other (Non-Grocery) Share							5.715	(0.308)
Warehouse Club Share							5.118	(0.181)
Number of Observations	65,852		65,852		65,852		65,852	
Adjusted R-squared	0.09		0.09		0.17		0.21	

Notes: The dependent variable is the household inventory-to-spending ratio (times 100). Maximum Age is the maximum age of household heads. Young Children is an indicator for whether children under the age of 6 are present in the house. White and Asian are indicator variables for the household head. College degree is an indicator for whether either household head has a college degree. Single Family Home is an indicator for whether the household lives in a single family home. ZIP Code House Prices are from Zillow. Income is the midpoint of the corresponding NielsenIQ bin. Perishable Share of Spending is the share spent on products with a time to expiry less than 0.58 months (just over two weeks). This cutoff is chosen so that ‘perishable’ here corresponds to perishability groups 1 and 2 in the model. The omitted store-type is grocery. Standard errors are in parentheses.

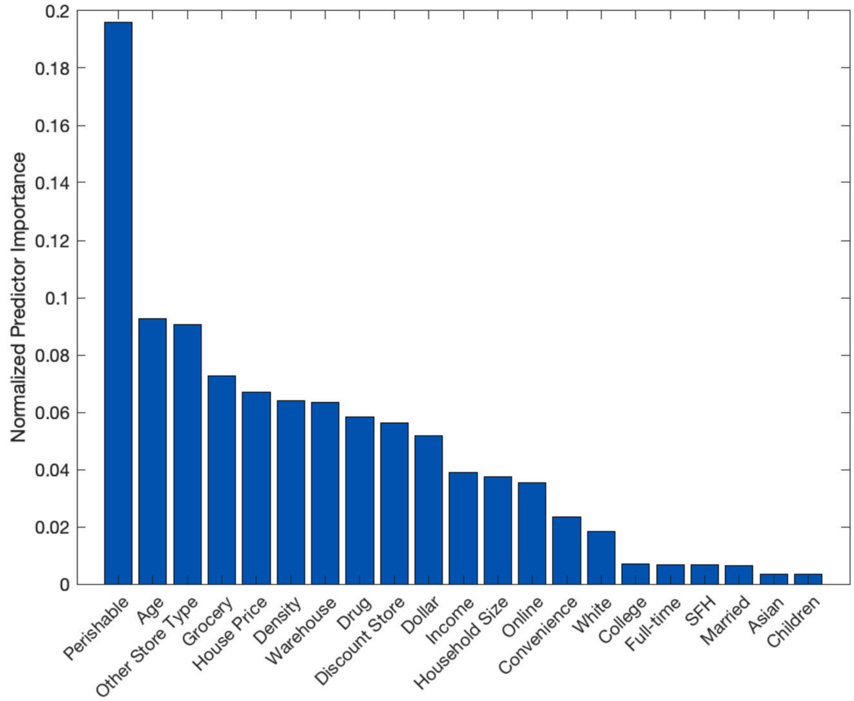

Fig. 4. Relative importance for inventory ratio prediction.
Notes: This figure shows the relative importance of each inventory ratio predictor. It measures the share of the reduction in mean-squared error due to each predictor. At each node where a predictor is chosen, the predictor's contribution is the difference between the MSE at the parent node and the average MSE of the child nodes (weighted by the number of observations going through each child node). The contribution is then summed over all nodes for which the predictor is chosen, weighted by the number of observations at each node as a share of the total sample size. We predict the inventory ratio using bootstrap aggregation (without random selection of predictors), 343 trees and a minimum leaf size of 27 observations.

Column 4 also includes the share of spending at different types of stores, with grocery store spending as the base. With the exception of discount stores (which feature similar types of products as grocery stores), spending more at non-grocery store types is associated with substantially higher inventory. There are a number of possible reasons for this. Shopping at warehouse clubs is likely indicative of a household's interest in buying in bulk and obtaining low prices. Store type may also be capturing some variation in perishability not picked up by our perishability measure. Finally, there may also be heterogeneity in pricing strategies across store types. We expect that households will stockpile more when shopping at retailers who offer large temporary discounts rather than everyday low prices. Including store type shares explains another 4 per cent of variation in the inventory ratio.

While most of these characteristics are statistically significantly correlated with inventory ratios, the overall explanatory power of even their combination is modest (R^2 of 0.21 or less). The fact that household characteristics explain a low portion of overall shopping behavior is consistent with findings reported in Hendel and Nevo (2006b), which documents that using household characteristics to predict the likelihood to purchase a product during a sale yields a low R^2 of under 0.03.

We might expect that conditioning on observed shopping choices will influence some of the coefficients on more fundamental characteristics. Interestingly, coefficients on age, education, household size, marital status, and labor force status do not change much. We see substantial changes for 'young children', 'income', 'density' and race. Households with young children consume more perishable products and both income and density are more correlated with choice of store type.

While we document correlations rather than causal relationships, in most cases the signs in Table 2 are consistent with theoretical predictions. Youth, full-time work, income, and young children are all be linked to the cost of time. We expect that households with a higher cost of time will prefer to save by stockpiling in response to deals observed while in the store, rather than by searching across stores, or shopping more frequently to take advantage of low prices on specific items. Education, income, and ZIP Code house prices are also positively correlated with wealth and therefore negatively related to financial constraints which would limit inventory accumulation.

ZIP Code population density and property type are likely to be correlated with storage space. The positive coefficient on single family home is consistent with this, but the coefficient on density is not. However, it is difficult to conclude much from the correlation with density given that many other factors linked to shopping behavior are also related to density. After conditioning on perishability and store choice, the coefficient on density becomes insignificant. In Appendix Table H.2, we use transaction data from a FinTech company to show that this effect may also be influenced by the extensive margin of food shopping. Conditional on income, households in denser locations tend to spend proportionally more on restaurants and less on grocery goods, which may prompt differences in the amount of accumulated inventory.

To further characterize households who hold high and low inventory as a share of spending, we also take a machine learning approach based on the variables in Column 4. This approach allows for non-linearities and interactions between variables and is similar to the approach we use in Section 3.2 to impute inventory holdings for SCF households.[9] Fig. 4 shows the relative importance of the different predictors (normalized to sum to one). It measures the share of the reduction in mean-squared

[9] We use the Matlab command fitrensemble with hyperparameter optimization to train the inventory prediction model. This method is similar to a random forest but requires all predictors to be used for every tree. For this application, we are more interested in understanding predictor importance and relationships than predicting inventory (though in any case alternative methods such as random forest and LSBoost have little effect on the quality of the predictions here). Using NielsenIQ weights, the optimized method uses 343 trees and a minimum leaf size of 27 observations.

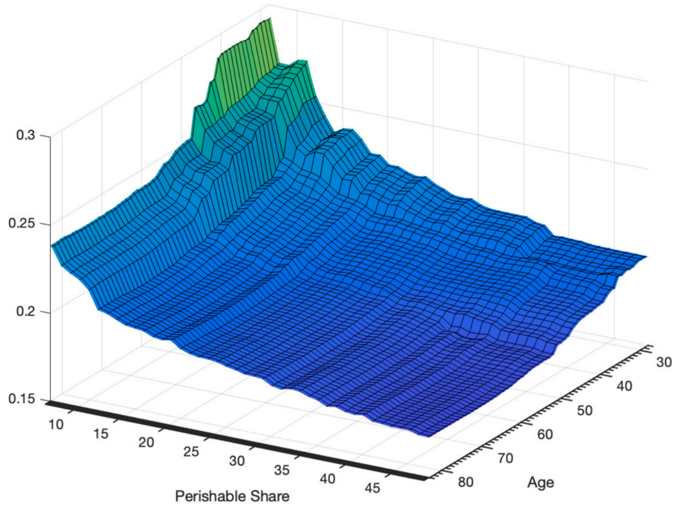

Fig. 5. Partial dependence profile: age and perishable share.
Notes: This figure shows the joint partial dependence profile for age and the share of purchases with a shelf-life of less than 0.58 months. To construct the partial dependence profile, model predictions are computed for each observation at counterfactual values of age and the perishable share (between the 2nd and 98th percentile), holding all other predictors fixed. The resulting predictions are then averaged over all observations in the dataset (applying NielsenIQ sample weights).

error due to each predictor. We find that the most important predictor is the perishable share followed by age.

Fig. 5 shows a joint partial dependence plot for the two most important predictors. The ratio of inventory to spending is increasing in age and declining in the perishable share. Fig. 6 shows partial dependence plots for the six most important predictors. These relationships are mostly monotonic and consistent with Table 2. The main exception is that inventory ratios are decreasing in population density for very low levels of density. This could be consistent with space constraints, or with a high trip fixed cost reflecting distance from the store. If we exclude shopping outcomes from the model, the most important variable is 'age', followed by 'density', 'house price', 'household size', 'white', and 'income'.

4. A model of optimal household inventory management

By setting aside working capital, households can reduce the average price they pay for consumer products. This can act as a substitute to the channel identified by previous work that has focused on taking more frequent shopping trips to take advantage of lower prices (Aguiar and Hurst 2007). In this way, people with a relatively high opportunity cost of time can obtain savings by stockpiling items when they are on sale instead of engaging in more frequent trips.

To understand the implications for borrowing behavior and portfolio allocation, we need to know the marginal financial return to allocating additional funds to household working capital. In this section, we use the NCP data to calibrate a model of optimal household inventory management. We then use the model to compute the net marginal return to household working capital investment, taking into account product depreciation costs and shopping trip fixed costs.

Our model builds on a previous literature, including our own work, which shows that consumers use stockpiling strategically to take advantage of temporarily low prices and to reduce the frequency of shopping trips. Much of that previous literature has exploited temporary price shocks to identify the stockpiling channel of consumer responses to these shocks. In contrast, we are interested in how allocating a marginal dollar to household working capital facilitates savings.

4.1. Model overview

In our continuous-time model, a household with an infinite horizon minimizes the cost of providing an exogenously given consumption flow subject to a working capital constraint. For simplicity, we assume that the flow of consumption is constant both between trips and across trips.[10] We define the return to working capital investment as the reduction in cost generated by relaxing the working capital constraint. There are two types of costs: a fixed trip cost and a variable cost per unit purchased. The variable cost reflects both the price charged by the store and the cost of storing the product between purchase and consumption. We assume that this storage cost corresponds to physical quality deterioration and calibrate it using shelf life data. The trip fixed cost (e.g., the opportunity cost of time spent shopping or pecuniary costs of travel) implies that even though the model is in continuous time, spending occurs at discrete, evenly-spaced dates endogenously chosen in advance by the household.

Households consume a continuum of goods i with varying perishability indexed by l. Each good is characterized by a rate of depreciation and a maximum shelf life beyond which it cannot be stored. Allowing for heterogeneity in perishability is important for matching the data. Perishable goods drive the frequent trips observed, while non-perishable goods are important for matching the substantial stockpiling observed in response to price changes (Hendel and Nevo 2006a; Baker et al. 2021).

We model the household's choice in two stages. First, we consider the household's in-store choice of how much of a product to purchase for storage given its observed price. This part of the model is broadly similar to Boizot et al. (2001) and Hendel and Nevo (2006a). If an item is not on sale, a purchase is only made if the household does not have enough inventory of that item to last until the next trip. If an item is on sale, the household replenishes inventory to a target level as in an (s, S) type model. We refer to this target as the stockpiling strategy. Our goal is to derive expressions for the average price the household pays per unit as a function of the stockpiling strategy, and the working capital required to facilitate the strategy. More stockpiling reduces the average price paid, but incurs depreciation costs and requires more working capital.

Second, given the first stage, we model the household's choice of the time interval between trips subject to a working capital constraint. This second stage is similar to standard inventory models, where a firm decides when to place orders in order to minimize costs of meeting demand (Arrow et al. 1951).[11] We refer to the first stage as "the stockpiling problem" and the second stage as "the trip-timing problem." To our knowledge, we are the first to integrate these two problems and incorporate a working capital constraint.

Finally, we explain how our model can in turn fit into a consumption and portfolio choice problem. This allows the working capital investment to be considered alongside traditional financial assets in a household's portfolio.

Our model incorporates two types of savings: buying items on sale ("deals") and buying in larger quantities ("bulk"). In turn, these drive two key relationships between unit prices and shopping trip frequency. Buying in bulk relates directly to the size of the trip (i.e., the amount spent per trip) and buying items on sale relates directly to the frequency of the trip (i.e., more frequent trips yield on average more items purchased on sale for a given trip size). Because buying large quantities

[10] This assumption can be relaxed. For the CES case, see Baker et al. (2021). Appendix B shows that for the type of products covered by the NielsenIQ data, this is a reasonable representation of consumer behavior.

[11] It is also similar to the steady state version of the model in Baker et al. (2021). An important difference is that Baker et al. (2021) captures intertemporal substitution behavior in response to an anticipated persistent consumption tax change; whereas here households take advantage of periodic sales and maintain a permanent base level of inventory.

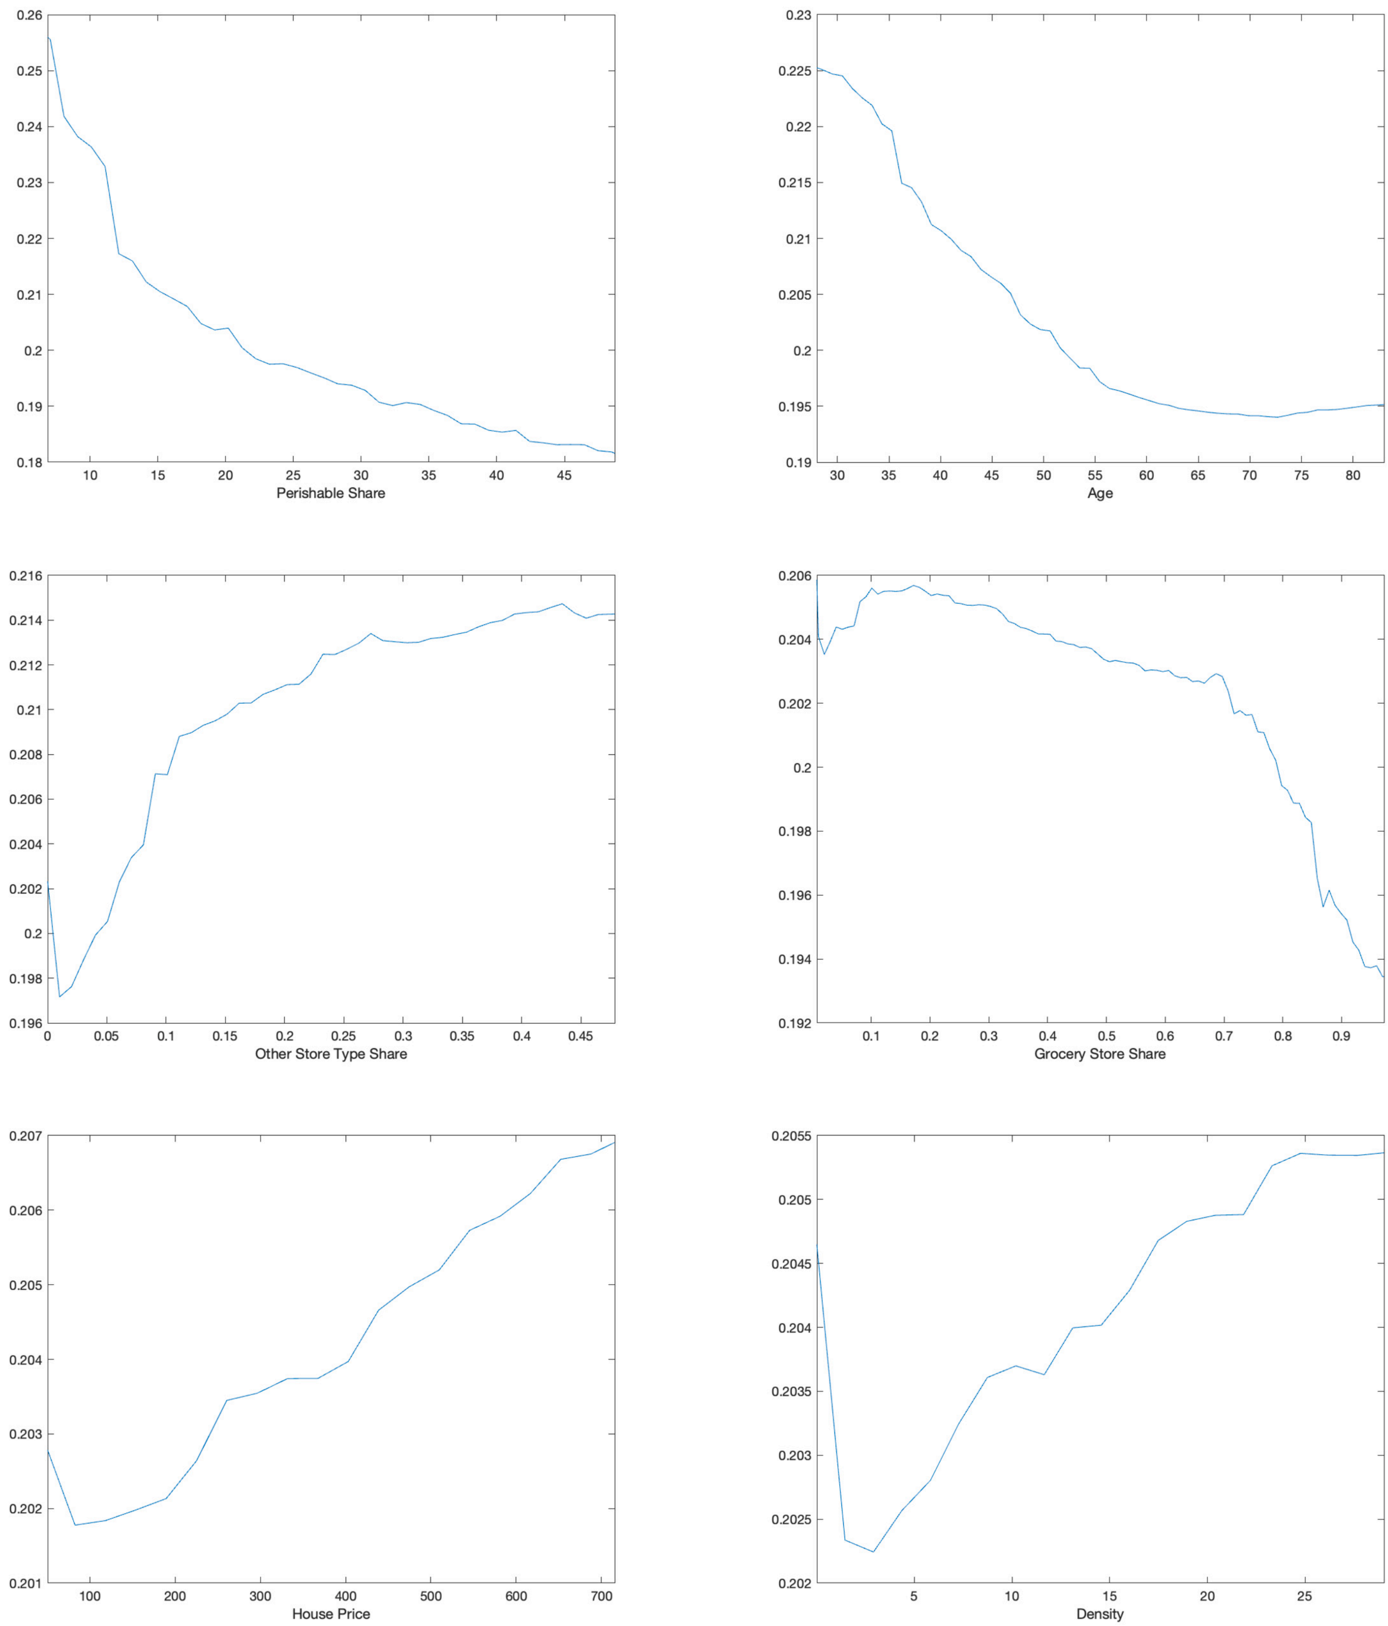

Fig. 6. Partial dependence profiles.

Notes: This figure shows partial dependence profiles for the six most important inventory ratio predictors. To construct the partial dependence profile, model predictions are computed for each observation at counterfactual values of the variable of interest (between the 2nd and 98th percentile), holding all other predictors fixed. The resulting predictions are then averaged over all observations in the dataset (applying NielsenIQ sample weights).

reduces trip frequency and the ability to take advantage of sales, there is a trade-off between the two types of shopping policies. Depending on various parameters (amount of household working capital, depreciation rate, shopping trip fixed cost, frequency and magnitude of sales, etc.), households may prefer one shopping policy over the other.

4.2. The stockpiling problem

We first consider the stockpiling problem from the perspective of a household who shops at trip interval Δ. The household chooses how much of each item to purchase for storage given the price observed on the current trip. In this section we characterize the relationship between the household's stockpiling choice and the price per unit. In Section 4.3, we integrate this relationship into the trip-timing problem and solve for the optimal trip interval and stockpiling choice jointly.

Although the stockpiling component of our model is similar in some respects to Boizot et al. (2001) and Hendel and Nevo (2006a), there are also some crucial differences. We assume that the household observes prices only upon entering the store and there is no further fixed cost of making purchases. Rather than focusing on a single product, we assume the household makes purchases to supply a consumption stream of a continuum of products. Sales rotate across products, but the share of products on sale is the same each trip. The decision to go to the store is not precipitated by a low price realization for a single product of interest so the trip interval Δ is constant. Our problem is effectively deterministic after aggregating across products.

For most nondurable consumer products and for purchase frequencies observed in our scanner data, prices tend to be at a modal level with temporary price discounts. Following previous studies (Boizot et al. 2001; Hendel and Nevo 2006a), we simplify retailers' price setting with two prices, the modal price, which we call full (or "list") price p_f, and the discounted (or "sale") price p_d. The discounted price is observed with probability x and the current price realization is assumed to be independent of past prices. We also assume price discounts are independent across products. Because the household consumes a continuum of products with prices independently drawn from the same distribution, the amount and quantity purchased on each trip and the inventory remaining at trip time are the same for all trips.[12]

We require the household to purchase an integer number of "packs" of each product. One pack is the quantity that must be purchased on the current trip in order to supply consumption until the next trip. We assume constant consumption of all products and therefore cannot have the household running out of a product between trips. The household's optimal stockpiling strategy takes an (s, S) form, where the boundaries depend on the current price p. The optimality of this strategy is proven for the continuous case by Hall and Rust (2007) and Sethi and Cheng (1997) and it is also shown there that the optimal policy rules $s^*(p)$ and $S^*(p)$ are decreasing in p.

Fig. 7a provides an example of the path of inventory of an individual retail product and the corresponding (s, S) policy rules for prices p_f and p_d. Intuitively, when the price is at its maximum level p_f, it does not make sense to buy more than is required for consumption before the next shopping trip because the price next trip cannot be any higher (and may be lower). Therefore, the household only makes a purchase at full price when it has no product left in stock. That is, the optimal policy at full price is $s^*(p_f)=0$ and $S^*(p_f)=1$ packs.

In contrast, when a product is on sale, the current price is lower than the *expected* price at the time of the next trip (i.e., $p_d < xp_d + (1-x)p_f \equiv E[p]$). Depending on depreciation costs, it may make sense for the household to buy more than one pack. Because p_d is the lowest possible price, the only aspect of the policy we still need to solve for is $s^*(p_d)$, i.e., $S^*(p_d)=s^*(p_d)+1$ because there is no fixed cost of purchasing once the household is already in the store. Hence, the target level at discounted price, $s_d \equiv s^*(p_d)$, fully summarizes the stockpiling strategy.

Let q denote the current quantity. The household's order quantity q^o when using strategy s_d is:

$$q^o = \begin{cases} s_d + 1 - q & \text{if } q \le s_d \text{ and } p = p_d, \\ 1 & \text{if } q = 0 \text{ and } p = p_f, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

When the deal price is observed, the household purchases one pack for immediate consumption until the next shopping trip and $s_d - q$ packs for storage. Note that s_d can equivalently be thought of as the number of trips the household is willing to buy in advance of consumption when the deal price is observed.

4.2.1. Expected price conditional on the stockpiling strategy

We would like to find an expression for the long-run quantity-weighted expected price paid conditional on the stockpiling strategy, $\bar{p}(s_d)$. To do this, we need to know the share of sale purchases for a given choice of s_d. Let d be the steady-state share of deal purchases. By setting a higher value of s_d, the household can raise deal share d for a given value of x. The long-run expected price paid conditional on strategy s_d is then:

$$\bar{p} = \frac{p_d E[q^o \mid p{=}p_d, q \le s_d]d + p_f E[q^o \mid p{=}p_f, q \le s_d](1-d)}{E[q^o \mid p{=}p_d, q \le s_d]d + E[q^o \mid p{=}p_f, q \le s_d](1-d)}. \quad (9)$$

That is, the numerator is the share of transactions occurring at p_d, times the average transaction value of orders at p_d, plus the share of transactions occurring at p_f, times the average transaction value of orders at p_f. The denominator is the expected number of packs purchased per transaction.

We know that conditional on p_f being observed and a purchase occurring ($q \le s_d$), exactly one pack will be purchased. Also, it will always be the case that $q \le s_d$ as long as q is below s_d at time zero. This gives us:

$$\bar{p} = \frac{p_d E[q^o \mid p{=}p_d]d + p_f\left(1-d\right)}{E[q^o \mid p{=}p_d]d + \left(1-d\right)}. \quad (10)$$

Next we will find expressions for d and $E[q^o \mid p{=}p_d]$.

Under the assumptions stated above, the price at which a transaction occurs follows a first-order Markov process. The rows of the transition matrix below correspond to the price at which transaction t occurs and the columns correspond to the price at which transaction $t+1$ occurs. The first row (column) of the transition matrix below corresponds to p_d and the second row (column) to p_f:

$$\Pi = \begin{pmatrix} 1-(1-x)^{s_d+1} & (1-x)^{s_d+1} \\ x & 1-x \end{pmatrix}. \quad (11)$$

The intuition is as follows. When a transaction has just occurred at p_d this means there are currently s_d+1 packs in stock. For the next transaction to occur at p_f it must be the case that no sale is observed for more than s_d trips in a row. As prices are i.i.d., the probability of this is $(1-x)^{s_d+1}$.

When a transaction has just occurred at p_f there is currently one pack in stock and a transaction must occur on the following trip. As prices are i.i.d., the probability the next transaction occurs at p_d is x and the probability it occurs at p_f is $1-x$.

Now that we have the transition matrix, solving for the steady state deal share d is straightforward (Appendix F provides intermediate steps):

$$d = \frac{x}{x + (1-x)^{s_d+1}}. \quad (12)$$

Next, we need to work out the steady-state average quantity purchased when a deal is observed, $E[q^o \mid p{=}p_d]$. The quantity purchased depends

[12] In Appendix G we extend this to the case of three prices with similar results. The three price case is considerably more complex (Boizot et al. 2001) and the computational cost of extending to a larger number of prices is substantial.

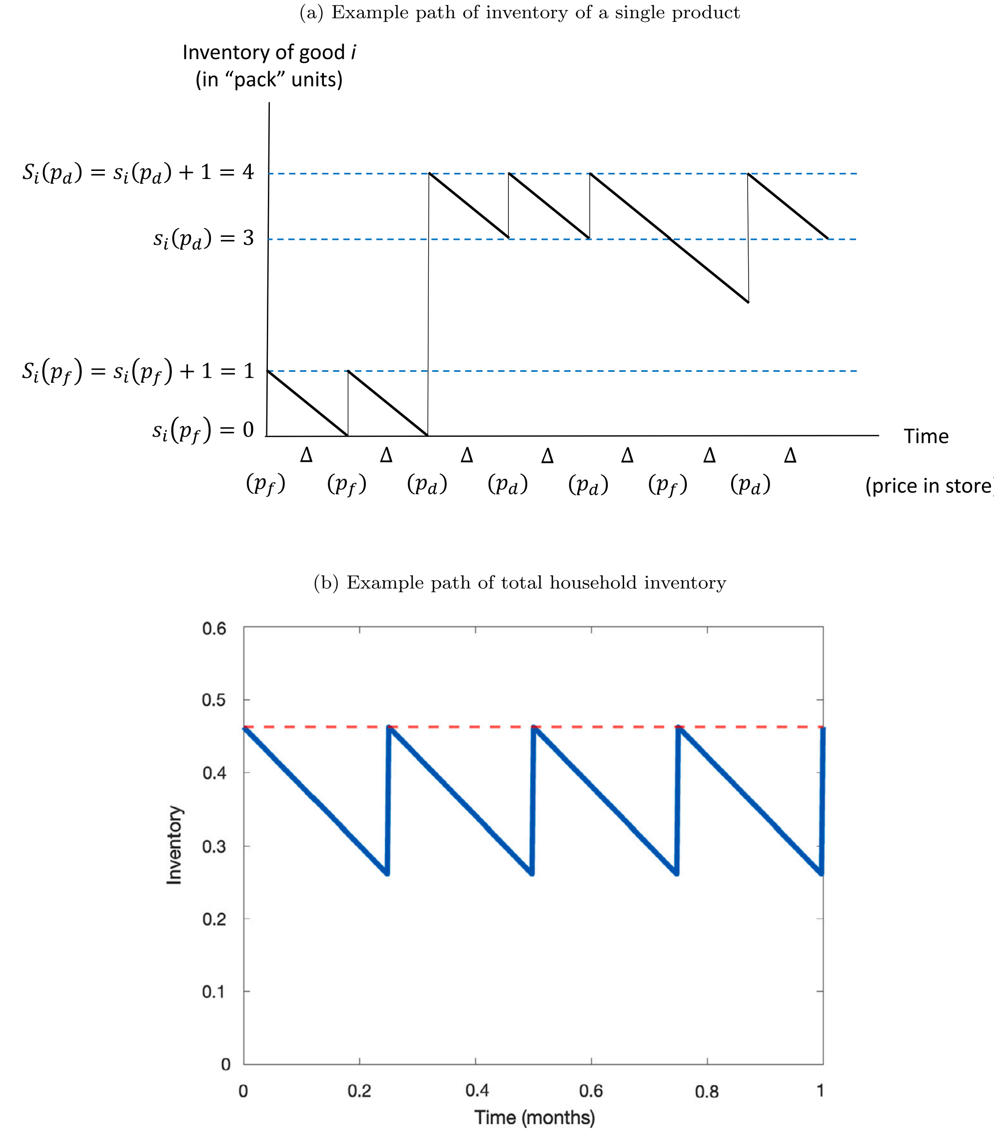

Fig. 7. Example paths of household inventory.

Notes: This figure shows two examples of the path of inventory. Panel (a) show the path of an individual retail product i. The blue dashed lines and the x axis display the (s, S) policies at full price p_f and discounted price p_d, respectively. In this example, $s_{d,i} = s_i(p_d) = 3$. Panel (b) shows an example of the path of inventory aggregated to the household level. In this example, the optimal trip interval is $\Delta = 0.25$ months, and optimal deal shopping strategies for product groups 1 to 4 are $s^*_{d,1} = s^*_{d,2} = s^*_{d,3} = 0$, and $s^*_{d,4} = 4$. The red dashed line displays total working capital, which is the sum of inventory and cash set aside to pay for the next shopping trip. Cash accumulates between trips at the rate at which inventory depletes, leaving total working capital constant.

on how long it has been since a deal was last observed. Regardless of when a deal was last observed, at least one pack is added (at a minimum, the pack consumed over the previous period is replaced). For each additional period when a sale is not observed, the household adds one extra pack upon next observing p_d. That is, if p_d is observed two trips in a row, the household buys one unit on the second trip. If p_d is next observed after two trips the household buys two units, and so on. The probability that t periods pass without a sale being observed is $(1-x)^t$. The expected quantity purchased conditional on a transaction occurring at p_d is therefore[13]:

$$E[q^o \mid p = p_d] = \sum_{t=0}^{s_d} (1-x)^t. \tag{13}$$

Substituting (12) and (13) into (10) and simplifying gives:

$$\overline{p} = \left(1 - (1-x)^{s_d+1}\right) p_d + (1-x)^{s_d+1} p_f. \tag{14}$$

(14) can also be obtained more directly by intuitive argument. The only case where the household pays p_f in the long-run is when there has been a sequence of $s_d + 1$ trips without sale (and only one pack is purchased at p_f). In all other cases the household pays p_d. The formal derivation is helpful, however, for two reasons. First, we have so far ignored depreciation costs. To accurately incorporate these costs it is necessary to keep track of the distribution of times between purchases. The distribution of time between purchase is also necessary for computing the level of inventory implied by stockpiling strategy $s(p)$. Second, the additional structure is helpful when we derive the solution to the considerably more complicated three price version of the model in Appendix G.

4.2.2. *Incorporating depreciation*

Standard inventory models, such as Arrow et al. (1951), typically include a storage cost per unit time which is expressed as a proportion of the total inventory. In our model, we directly account for product de-

[13] We also verify (13) using a simulation for both small and large values of s_d.

terioration as part of the individual product stockpiling problem before aggregating.

There are two equivalent ways to think about the depreciation cost. We can convert the price paid in store to an "effective price," which reflects quality deterioration over the holding period. We can also think of the depreciation cost as modifying the quantity that must be purchased to meet consumption needs: if items randomly go bad, more needs to be purchased initially (with the additional amount required being determined by the holding period). As the total amount spent on good i each trip is $p_i q_i^o$, in practice it does not matter whether we apply the depreciation cost factor to the price or the quantity.

If a product deteriorates exponentially at rate δ, a unit purchased at time zero for consumption following trip t effectively costs $e^{\delta t \Delta}$ times more than a product purchased for immediate consumption (recall that Δ is the time between trips, so $t \cdot \Delta$ is the total time between purchase and consumption). To adjust the price function (14) to account for these costs, we need to consider the distribution of inter-purchase times. This is because the level of inventory prior to a purchase determines the holding period for the items purchased. For example, if a sale is observed every trip, the household will have s_d packs in stock immediately prior to each trip, and purchase one pack each trip. Under a first-in first-out approach, each pack would be opened s_d trips after it was purchased. The depreciation factor associated with each pack would therefore be $e^{\delta \cdot s_d \cdot \Delta}$. In contrast, if inventory is zero immediately prior to a trip where p_d is observed, the first pack purchased will be opened immediately (and have a depreciation factor of one), the second pack will be opened after the next trip, and so on.

It is convenient to incorporate depreciation costs into our expression for $E[q^o \mid p = p_d]$. Recall that for each additional trip when a sale is not observed, the household adds an extra pack upon next observing p_d. Under a first-in first-out approach, the incremental pack is held for $s_d - t$ trips before consumption, where t is the number of trips without a sale. Therefore, depreciation inflates the cost of the incremental pack by a factor of $e^{\delta(s_d - t)\Delta}$. This gives:

$$E[q^o \mid \mathrm{p{=}p_d}] = \sum_{t=0}^{s_d} e^{\delta(s_d - t)\Delta}(1-x)^t. \tag{15}$$

As we increase the number of trips without a sale, t, the depreciation cost of the incremental pack $e^{\delta(s_d - t)\Delta}$ declines because that pack will not be held for as long prior to consumption. The expression for the average price after incorporating depreciation is therefore:

$$\overline{p}(s_d, \Delta) = p_d x \sum_{t=0}^{s_d} e^{\delta(s_d - t)\Delta}(1-x)^t + p_f (1-x)^{s_d + 1}. \tag{16}$$

Because we aggregate over a continuum of goods indexed by $i \in [0,1]$, assuming that sales are i.i.d., $\overline{p}(s_d, \Delta)$ is the quantity-weighted average price per unit on every trip.

In addition to exponential quality deterioration, we also incorporate a product shelf life constraint, $\bar{t}$. This is an upper bound on the time (measured in months) a product can be stored:

$$(s_d + 1)\Delta \leq \bar{t}. \tag{17}$$

This is intended to make the model more realistic. While many products do deteriorate gradually from the date of purchase, there is also typically a point at which quality falls below the minimum required for consumption. The expiration date of the product provides guidance to the consumer regarding when this point has been reached, and this is ultimately how we will calibrate $\bar{t}$. The method for determining expiration dates depends on the product. For some products, it is a question of food safety. For others, expiration dates are driven by perceived quality falling below a cutoff. Regardless of the reason, we think it is unlikely that households would set s_d so high that they would not consume the item prior to expiration.

4.2.3. Aggregate inventory

The amount of working capital required to facilitate strategy s_d is equal to the (deterministic) maximum value of inventory over the trip cycle. Inventory is at its maximum immediately following a trip. Therefore, to incorporate a working capital constraint into the final problem we want an expression for the value of inventory immediately following a trip. We will then gradually relax the constraint to compute the marginal net return to household working capital.

Intuitively, the higher the value of s_d, the higher the level of inventory will be when going to the store. If the household does not stockpile at all (i.e., $q^o = 1$), inventory is exactly zero at the time of the next trip. We first derive an expression for the inventory immediately prior to a trip, the number of packs in stock $I(s_d)$. As these packs were all purchased on sale, they cost p_d. To get the value of inventory following each trip, we add the value of purchases made on each trip. This is discussed in Section 4.3. Here, we focus on $I(s_d)$.

To obtain total inventory, we aggregate over the continuum of goods indexed by $i \in [0,1]$. In general, the household has $s_d - t$ packs in stock of share $x(1-x)^t$ of products immediately prior to each trip, where t is the number of trips since the last sale. Integrating across all products i, the total number of packs in stock immediately prior to every trip is:

$$I(s_d) = 1_{\{s_d > 0\}} x \sum_{t=0}^{s_d - 1} (s_d - t)(1-x)^t. \tag{18}$$

4.3. The trip-timing problem

Next, we integrate the price equation (16), shelf life constraint (17) and the inventory equation (18) into the trip-timing problem. The household will choose both s_d and trip interval Δ to minimize the cost of providing an exogenous consumption stream, subject to a working capital constraint. The cost per trip can be decomposed into two components: a fixed cost and a variable cost which depends on the quantity of products purchased.

We divide products into perishability groups indexed by l. Constant continuous consumption of perishability group l is $C_l(t) = C_l$. We use k to denote the fixed trip cost and P_l to denote the per unit price of group l. P_l will depend on the trip interval Δ and also on the household's stockpiling strategy, $s_{l,d}$. We allow the household to make a separate stockpiling choice for each perishability group l. So far we abstracted from this for simplicity. Each perishability group still contains a continuum of individual products i.

Each trip, the household must purchase enough of each good to last until the next trip, and we previously defined a pack as containing exactly this amount. We do not allow households to set different values of Δ for different goods. Although setting different values of Δ allows households to reduce depreciation costs, this is more than offset by the increase in trip fixed costs associated with maintaining multiple trip schedules. For more detail on this tradeoff, see Bartmann and Beckmann (1992).

We then work out the quantity per pack. For a good in group l that deteriorates exponentially at a rate δ_l, the quantity that must be purchased to satisfy continuous consumption flow C_l over the trip interval Δ is:

$$Q_l(\Delta) = \int_0^{\Delta} e^{\delta_l t} C_l \, dt = \begin{cases} (e^{\delta_l \Delta} - 1)\frac{C_l}{\delta_l} & \text{if } \delta_l > 0, \\ C_l \Delta & \text{if } \delta_l = 0. \end{cases} \tag{19}$$

That is, when the time between trips is Δ and the household buys quantity $Q_l(\Delta)$ each trip, inventory next hits zero precisely when the next shopping trip is scheduled to occur.[14]

[14] While depreciation between trips (i.e., while the pack is in storage) is reflected in the effective price (16) as part of the stockpiling problem, depreciation within trips (i.e., as the pack is being consumed) is reflected in $Q_l(\Delta)$.

Table 3
Calibration of model parameters.

Name		Perishability Group (l) 1	2	3	4	Source/target
Depreciation rate	δ_l	4.29	1.76	0.34	0.00	See Appendix C
Shelf life (in months)	$\bar{t}_l$	Δ^*	0.39	2.06	27.92	
Consumption flow (in %)	C_l	12.47	11.52	12.90	63.1	Match NCP expenditure shares.
Deal probability	x_l	0.21	0.30	0.29	0.27	Match NRP price moments.
Full price	$p_{l,f}$	1.07	1.07	1.08	1.08	Match NRP price moments and $E[p_l] = 1$.
Deal price	$p_{l,d}$	0.74	0.83	0.80	0.78	
Regression coefficients	$\hat{a}_{0,l}$	0.77	0.82	0.84	0.81	Estimation of (26) by WLS.
	$\hat{a}_{1,l}$	3.00	0.94	4.41	0.77	
Bulk discount function	α_l	0.78	0.81	0.82	0.77	α_l, β_l match relation between NCP pack size and unit price.
	β_l	3.05	0.92	4.30	0.73	
	$\hat{\sigma}_l$	2.63	1.56	3.18	1.16	σ_l maximizes within-R^2 of (26).
Trip fixed cost	k	0.0139				Match NCP trip interval.

Notes: This table shows the model parameters by perishability group. Group 1 contains the most perishable products and Group 4 contains the least perishable. The calibration approach is described in Section 4.4. In the model, one unit of a product from group l stored for a period of t months since purchase provides consumption of $e^{-\delta_l t}$ units if $t < \bar{t}_l$ and 0 units if $t \geq \bar{t}_l$. C_l is the % of total NielsenIQ spending accounted for by each group. x_l is the probability of a sale, $p_{l,d}$ is the price in the event of a sale, and $p_{l,f}$ is the full price. $\hat{a}_{0,l}$ and $\hat{a}_{1,l}$ are weighed least squares coefficients of regression equation (26), estimated separately for each group. In the model, we normalize the price of the standard pack size to one, and the price of other pack sizes reflect percentage deviations from the standard pack size. We therefore set $\alpha_l = \frac{\hat{a}_{0,l}}{\hat{a}_{0,l}+\hat{a}_{1,l}e^{-\hat{\sigma}_l}}$ and $\beta_l = \frac{\hat{a}_{1,l}}{\hat{a}_{0,l}+\hat{a}_{1,l}e_l^{\hat{\sigma}}}$. $\hat{\sigma}_l$ is chosen to maximize the within-R^2 of (26). Note, because the implied shelf life for group 1, $\bar{t}_1$, is somewhat lower than the data trip interval (which we calibrate k to match), we set $\bar{t}_1 = \Delta^*$, the household's optimal trip interval. This effectively means that products from group 1 cannot be stockpiled.

In addition to stockpiling in response to temporary deals, households can also save by buying larger pack sizes and paying a lower per unit price. We incorporate bulk discounts by multiplying the expected price function (16) by a bulk discount function $b(Q_l)$:

$$P_l(\Delta, s_{l,d}) = b(Q_l(\Delta)) \cdot \bar{p}_l(s_{l,d}, \Delta), \tag{20}$$

where $b(Q_l)>0$, $b'(Q_l)<0$ and $b''(Q_l)>0$. (We discuss the calibration of $b(Q_l)$ in Section 4.4.4.) Note that the expected price is achieved with certainty and $P_l(\Delta, s_{l,d})$ is therefore deterministic. Empirically, larger trip sizes correspond to a household either consuming more or shopping less frequently.

Given that trips are evenly spaced with endogenous trip interval Δ, the average cost of providing the exogenous consumption flow is:

$$\frac{k + \sum_l P_l(\Delta, s_{l,d}) \cdot Q_l(\Delta)}{\Delta}. \tag{21}$$

Next, we need to incorporate the working capital constraint. (18) provides $I_l(s_{l,d})$, the number of packs in stock of perishability group l immediately prior to a trip. The amount of working capital required to facilitate a given set of stockpiling strategies $\{s_{l,d}\}_l$ and trip interval Δ is the maximum inventory over the trip cycle (summed over all groups l), which occurs immediately following a trip. At this point in time 100% of household working capital is held as stored inventory goods. Between trips, the inventory share of working capital gradually declines through consumption and depreciation and is replaced with accumulated cash used to pay for the next trip. The paths of aggregate inventory and household working capital are illustrated in Fig. 7b.

As all stockpiling occurs at price $p_{l,d}$, the inventory of group l prior to a trip is $I_l(s_{l,d})Q_l(\Delta)p_{l,d}$. This is the number of packs in stock, times the pack size in number of consumption units, times the price per consumption unit. Note that we allow the price distribution to vary across perishability groups. To compute the working capital required to support $(\Delta, \{s_{l,d}\}_l)$, we need to add the value of products purchased on a single trip, $P_l(\Delta, s_{l,d})Q_l(\Delta)$. The maximum allowable level of working capital is exogenous, denoted by $\bar{I}$, and we will compute the marginal return to working capital by gradually increasing $\bar{I}$:

$$\sum_l \left[P_l(\Delta, s_{l,d})Q_l(\Delta) + I_l(s_{l,d})Q_l(\Delta)p_{l,d} \right] \leq \bar{I}. \tag{22}$$

We now have all the elements we need. The household minimizes the average cost (21) of providing the exogenous consumption flow C_l subject to working capital constraint (22) and a restriction on storage time for each perishability level l, i.e., shelf life constraints (17):

$$V(\bar{I}) = \min_{\Delta, \{s_{l,d}\}_l} \frac{k + \sum_l P_l(\Delta, s_{l,d})Q_l(\Delta)}{\Delta} \tag{23}$$
$$s.t. \quad \sum_l \left[P_l(\Delta, s_{l,d})Q_l(\Delta) + I_l(s_{l,d})Q_l(\Delta)p_{l,d} \right] \leq \bar{I},$$
$$(s_{l,d} + 1)\Delta \leq \bar{t}_l \ \forall\, l.$$

Ultimately, we are interested in the relationship between the dollar amount invested in household working capital, $\bar{I}$, and cost of providing the consumption stream. In order for a particular shopping strategy to be feasible, the level of inventory immediately following a trip must not exceed the amount of household working capital $\bar{I}$. We will solve the problem for different levels of $\bar{I}$, and use this to compute the return to investing in household working capital (i.e., marginally increasing $\bar{I}$). The investment payoff will be the reduction in the cost V, so that we can define the marginal (net) return to household inventory management:

$$\text{Marginal (net) return: } r_{\bar{I}}(\bar{I}) = -V'(\bar{I}). \tag{24}$$

4.4. Calibration

We calibrate the model by choosing parameters to match a number of data moments, summarized in Table 3. To solve the model, we define a grid over trip intervals Δ and bargain-hunting strategies $\{s_{l,d}\}_l$. We then search over all combinations for which the household working capital constraint and shelf life constraints are satisfied and find the combination that minimizes the cost function.

4.4.1. Time units, consumption flows, and expenditure shares

Trip length Δ is expressed in months. For example, Δ=0.23 implies a trip interval of 0.23 months (i.e., one week). To assist with calibrating the bulk discount function, below we define a standard trip size $\hat{Q}_l$ as the trip size corresponding to the optimal choice of Δ in the absence of bulk discounts. Problem (23) then scales with consumption if we assume that the trip fixed cost k and standard trip size $\widehat{Q}_l$ both scale with consumption. We therefore set $\sum_l C_l = 1$, which—because we express the trip interval in monthly time units—implies that the household consumes one unit of total consumption over the course of one month, and express the trip cost and standard trip size as shares of monthly total consumption. Because $E[p_l] = 1$ and $\sum_l C_l = 1$, C_l is also the expenditure share, and we calibrate it to the observed expenditure shares of perishability group l in the NCP.

4.4.2. Depreciation and shelf life

Depreciation costs δ_l are monthly rates with continuous compounding. For example, if $\delta_l = 0.5$, each unit purchased corresponds to $e^{-0.5} = 0.61$ units one month later. To calibrate these costs, we start by manually assigning a shelf life for each NielsenIQ product module using information on product life from FSIS. We then allocate each NielsenIQ product module to one of four perishability groups. For each perishability group l, we will calibrate a depreciation rate δ_l and a shelf life $\bar{t}_l$.

We rank NielsenIQ product modules by shelf life and then determine cutoffs for perishability groups. The precise cutoff point becomes less important as products become more storable. This is because the marginal price savings from additional stockpiling are extremely low once products are being bought several months in advance. The precise shelf life is therefore unimportant for products that can be stored for several months with only a negligible decline in quality. We define the most storable group ($l = 4$) as NielsenIQ product modules with a shelf life greater than 6 months. At 6 months, over 99% of the possible price reduction through stockpiling has already been exploited.

In contrast, variation in shelf life is important for more perishable items. We therefore try to choose the cutoffs for the remaining groups sensibly so that we can capture the variation well with a small number of groups. We describe the procedure in Appendix C.

Table 3 shows that the least perishable group is by far the largest, accounting for 63% of NCP expenditure. We want to ensure that $\bar{t}_l$ is at least as long as the data trip interval for all groups. We therefore set $\bar{t}_1 = \Delta$, which implies that $s_1^* = 0$ (i.e., the household buys exactly the amount it requires to supply consumption until the next trip). This is a fairly minor adjustment as the average shelf life in perishability group 1 is 0.16 months, compared with a data trip interval of 0.28. In practice, the exact value of $\bar{t}_1$ is in any case uncertain. Some products may be stored for a bit longer than the standard shelf life, or consumed close to the trip time rather than continuously.[15]

4.4.3. Price process

We primarily use the NRP to calibrate the price process. To calibrate x, p_d, and p_f, we ideally want to match the distribution of posted prices. In the model, we normalize $E[p] = xp_d + (1-x)p_f = 1$. This implies a value for p_f given p_d and x and means we need two moments to calibrate the price process. We calibrate the price ratio $\frac{p_f}{p_d}$ and the deal frequency x jointly to match the skewness and relative variance (variance divided by the mean) of NRP prices. We do this separately for each perishability group l. Intuitively, the more negative the skewness, the smaller is x. Under our assumption of two price points, a symmetric distribution corresponds to $x = 0.5$. Negative skewness corresponds to $x < 0.5$. Higher relative variance corresponds to a higher price ratio (i.e., larger discounts).

We compute the price moments using 2013 and 2014 data. Stores report prices weekly and we focus only on UPC-store combinations for which units are sold in at least 49 weeks of the year. This ensures that products in the sample are sold in every month and there are enough observations to estimate moments at the UPC-store level. There are around 81 million UPC-store combinations in the NRP satisfying this criterion. For each UPC-store combination, we compute the mean, variance, and skewness of prices. Even after conditioning on UPCs the store stocks all year round, there are still occasionally some anomalies in reported prices (some of these are likely errors). To address concerns about these outliers, we compute the ratio of maximum to minimum prices for each UPC-store-year and drop cases where the ratio is greater than 5.

We compute moments at the UPC-store level to capture temporary sales for the same product at the same store over time. We do not want to incorporate differences in average prices across stores. Ultimately, we need to calibrate price distributions for the four model product groups, meaning that we need to aggregate. A simple approach would be to compute average normalized variance and skewness for each group, but some UPC-store combinations have very low normalized variances and it is not clear that we want skewness from these cases to contribute to average skewness. We select only the middle quintile of the store variance distribution for each UPC and then compute average variance and skewness within this set of stores for each UPC. We repeat this for UPCs within the same product module to come up with representative moments for each product module. We then weight each product module by its NCP spending share and compute the weighted average variance and skewness for each of the four model groups.

Finally, we define a grid over the discount probability x_l and the discount size $\frac{p_{l,f}}{p_{l,d}}$ and compute the implied proportional deviation in variance and skewness relative to the data. We then select the combination of $(x_l, \frac{p_{l,f}}{p_{l,d}})$ that minimizes the root mean squared error.[16]

4.4.4. Bulk discount function

We specify the bulk price discount function b to match bulk discounts observed in NRP data using the following functional form:

$$b(Q_l) = \alpha + \beta e^{-\sigma \frac{Q_l}{\widehat{Q}_l}}. \tag{25}$$

Unit prices decline as the quantity purchased per trip Q_l increases. $\widehat{Q}_l$ is the trip size associated with purchasing standard packs of each item and we will calibrate parameters (α, β, σ) such that $b(\widehat{Q}_l) = 1$. In the NRP, we define the "standard pack size" as the second quintile of the pack size distribution for each product. In the model, $\widehat{Q}_l = Q_l(\widehat{\Delta}) = \frac{C_l}{\delta_l}(e^{\delta_l \widehat{\Delta}} - 1)$, where $\widehat{\Delta}$ is the optimal trip interval in the model without bulk discounts (i.e., the optimal trip interval assuming $b(Q_l) = 1 \; \forall \, Q_l$).

Fig. 8 shows that the calibrated function matches the data well in several respects: unit prices decay exponentially with pack size and converge to some level above zero. As pack sizes become very small, unit prices increase but do not become arbitrarily large. We normalize the effective price per unit in the absence of discounts to $P_l(\widehat{\Delta}, 0) = 1$ (i.e., with $s_{l,d} = 0$ and $b(Q_l) = 1 \; \forall \, Q_l$). The parameter α is interpreted as one minus the maximum % savings that can be obtained from buying in bulk. Appendix C describes the calibration procedure in more detail.

We compute total expenditure for each product and use this to weight our regressions. We calibrate the parameters of the function $b(Q_l)$ by estimating a and b of the following relationship with weighted least squares separately for different values of σ:

$$\text{Unit Price}_{p,q} = a + be^{-\sigma \text{Units}_{p,q}}. \tag{26}$$

[15] This type of behavior is a realistic violation of the constant consumption assumption we make in Section 3. However, as we discuss in Appendix B, violations similar to this one (i.e., consuming perishable products early in the trip cycle) would not have a substantial effect on the inventory calculation.

[16] Formally, we minimize the root mean squared deviation in variance and skewness relative to the data, where hats denote data moments:

$$RMSE = \sqrt{\frac{1}{2}\left(\frac{\text{Var} - \widehat{\text{Var}}}{\widehat{\text{Var}}}\right)^2 + \frac{1}{2}\left(\frac{\text{Skew} - \widehat{\text{Skew}}}{\widehat{\text{Skew}}}\right)^2}.$$

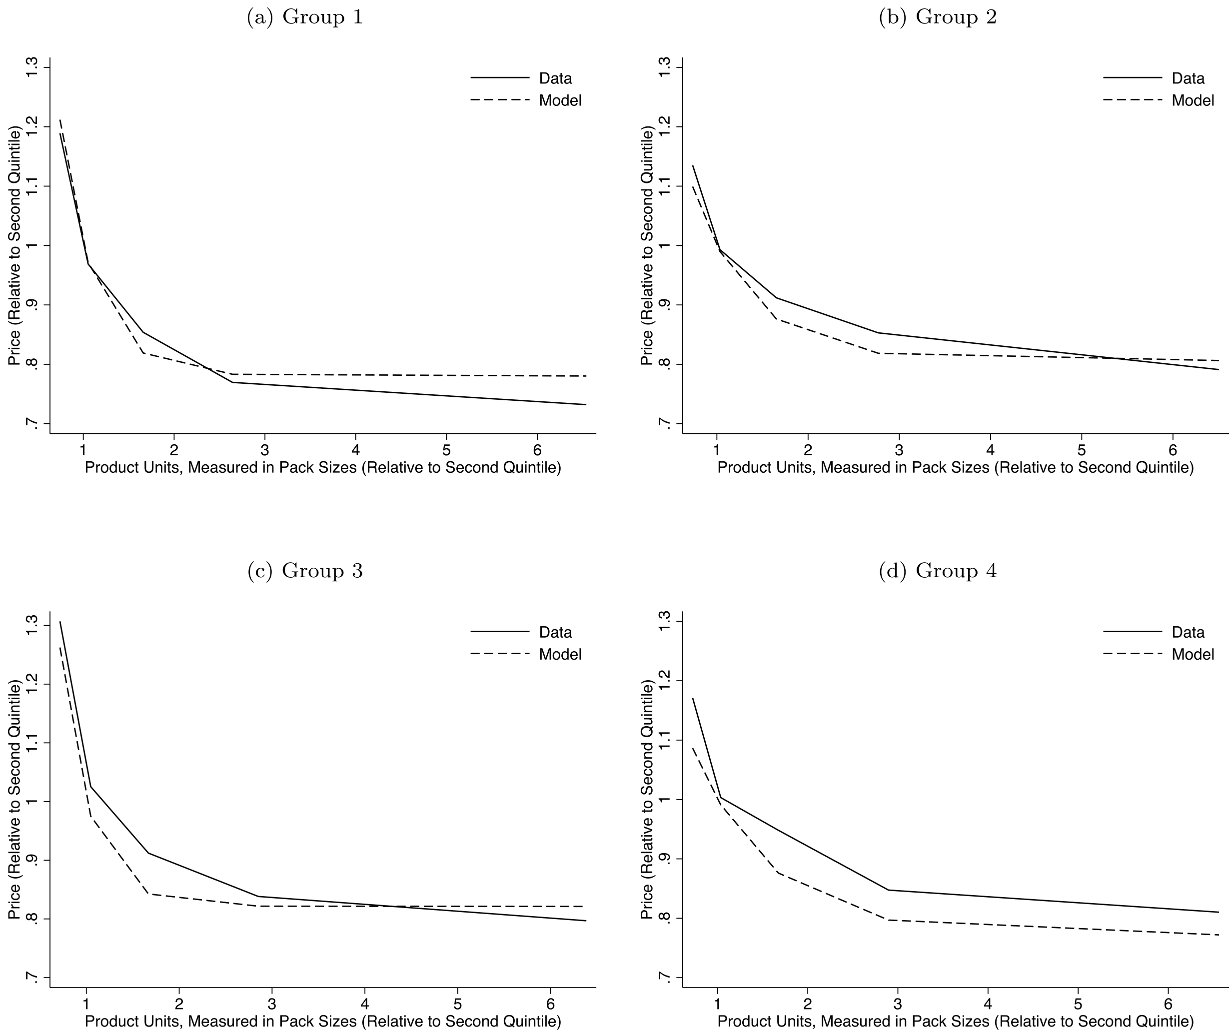

Fig. 8. Bulk calibration.
Notes: We first compute pack size quintiles for each product-ZIP3 combination. We then compute the median number of units and unit price for each product-pack size quintile-ZIP3, weighted by expenditure. We normalize both prices and units by dividing by the second quintile price and units. The range of available sizes varies substantially across products. As we want to measure the savings obtained by increasing pack size uniformly across all products, we ensure that all products have a common range of normalized units. This means the set of products does not change along the x-axis. To achieve this, we create a number of pack size bins over the range 0.5 to 10 (the cutoffs are 0.5, 1, 1.5, 2, 5, and 10). For products where price and units are missing for a particular bin, we impute price using the unit price in the closest bin. We impute units as the weighted average (normalized) units in the bin across all products. We then estimate (26). The dashed line shows the relative price we assume in the model: $\text{Price} = \alpha + \beta e^{-\hat{\sigma}\text{Units}}$ and Units is the weighted average normalized units in units bin q across all products. The solid line is constructed by computing the weighted average normalized retail price in units bin q across all products.

$\text{Unit Price}_{p,q}$ is the standardized unit price of product p at pack size q and $\text{Units}_{p,q}$ is the standardized number of units of product p at pack size q. We then choose $\hat{\sigma}$ to maximize the within-R^2. We perform this procedure separately for each level of perishability l. In the model, we normalize the price of the standard pack size to one, and the price of other pack sizes reflects percentage deviations from the standard pack size. We therefore calibrate α and β using $\alpha = \frac{\hat{a}}{\hat{a}+\hat{b}e^{-\hat{\sigma}}}$ and $\beta = \frac{\hat{b}}{\hat{a}+\hat{b}e^{-\hat{\sigma}}}$.

4.4.5. Trip fixed cost

We set the fixed cost per shopping trip so that the trip interval in the model matches the average interval between grocery trips (a household's modal shopping channel that makes up more than one third of their total NielsenIQ spending) in the NielsenIQ data. The average trip interval across households is 0.28 months (or slightly more than one week) and the corresponding fixed cost per trip is 1.39% of monthly consumption. This calibration accords well with the fixed shopping cost of \$4.85 estimated in Baker et al. (2021) using shopping responses to sales tax changes. Given average monthly NielsenIQ spending of approximately \$375, \$4.85 would be equivalent to a trip cost of $k = 1.29\%$ of monthly consumption.

4.5. Implications for portfolio choice

In the model, the working capital constraint and consumption are exogenous. The model should therefore be considered as one component of a higher-level problem in which the household chooses consumption and allocates assets to several investments, including working capital. For example, to understand the implications of the household-specific returns from inventory management for households' participation in risky financial markets, we briefly sketch out how our model fits into a static portfolio choice problem.

We consider the effect of working capital on the cost of supplying consumption to be analogous to interest earned on an investment. The household has access to three investment opportunities: working capital; a risk-free bond; and a risky asset, which could be thought of as the

Table 4
Financial returns to household inventory investment.

Work. Cap.	Inventory	Cash	% Savings:		Interval Δ^*	$s^*_{4,d}$	Return (%):	
($\bar{I}$, % cons.)	(% spend.)	(% $\bar{I}$)	Deal	Bulk	(months)	(trips)	Marg.	Avg.
2.5	1.8	35.1	6.2	12.8	0.23	2	215.2	
5.0	4.7	19.0	9.9	15.7	0.26	5	56.6	215.2
7.5	7.4	13.6	10.8	17.3	0.28	7	28.6	135.9
10.0	10.7	9.4	11.8	16.1	0.27	10	13.7	100.1
12.5	13.6	7.8	11.9	16.7	0.28	12	6.3	78.5
15.0	16.5	6.6	12.0	17.1	0.28	14	3.0	64.1
17.5	19.4	5.7	12.1	17.2	0.28	16	1.6	53.9
20.0	22.1	5.0	12.1	17.2	0.28	18	0.8	46.4
22.5	24.9	4.4	12.1	17.2	0.28	20	0.5	40.7
25.0	27.6	4.0	12.2	17.2	0.28	22	0.2	36.3
27.5	30.3	3.6	12.2	17.2	0.28	24	0.1	32.7
30.0	34.3	3.3	12.2	17.2	0.28	27	0.1	29.7
32.5	37.2	3.1	12.2	17.2	0.28	29	0.0	27.2
35.0	39.9	2.9	12.2	17.2	0.28	31	0.0	25.1

Notes: This table is constructed by solving the model for different values of the working capital constraint $\bar{I}$, increasing it by 2.5% of (exogenous) annual consumption in each row. The working capital ratio in column 1 is available working capital $\bar{I}$, expressed as a percentage of annual consumption. Column 2 shows the value of inventory immediately prior to a trip as a percentage of total annual spending. Column 3 shows the annual average value of cash set aside to pay for the next trip as a share of annual total working capital. The remaining share of working capital is invested in inventory. Columns 4 and 5 show in-store savings achieved in % of base spending, which is annual spending assuming no stockpiling ('untargeted' or 'inattentive' shopping, $s_{l,d} = 0$) and the trip interval is the interval associated with purchasing the standard pack size of each product. Deal savings are in-store savings due to buying an item on sale. Bulk savings are in-store savings due to buying a larger pack size. Column 6 shows the length of time of the optimal interval Δ^* between trips, measured in months. $s^*_{4,d}$ in column 7 is the optimal deal shopping strategy for goods with a shelf life of at least six months (group 4), expressed in the number of trips the household is willing to purchases the product in advance of consuming it when the product is on sale. The financial returns in columns 8 and 9 incorporate not only in-store savings but also depreciation and trip fixed costs. The average return is computed relative to a working capital benchmark of 2.5% of annual consumption.

market portfolio. It maximizes expected utility of end-of-period wealth by solving the following problem:

$$\max_{\lambda_{\bar{I}},\lambda_f,\lambda_m} EU\left((1+\tilde{r}_p)w\right) \tag{27}$$

$$s.t. \quad \tilde{r}_p = \frac{1}{w}\int_0^{\lambda_{\bar{I}} w} r_{\bar{I}}(x)dx + \lambda_f r_f + \lambda_m \tilde{r}_m, \tag{28}$$

$$1 = \lambda_{\bar{I}} + \lambda_f + \lambda_m. \tag{29}$$

$\lambda_{\bar{I}}$ is the share of initial wealth w allocated to working capital (so $\bar{I} = \lambda_{\bar{I}} w$) with marginal return $r_{\bar{I}}$, λ_f is the share allocated to the risk-free bond with return r_f, and λ_m is the share allocated to the risky asset with stochastic return $\tilde{r}_m$ (e.g., stock market).

$r_{\bar{I}}$ is the working capital return function we solve for using our model. While the risk-free bond and risky asset returns do not depend on the amount invested, the return on working capital depends on the amount invested. Consistent with our model assumptions, we treat the working capital investment as a risk-free asset.

Assuming consumers are risk averse, they choose $\lambda_{\bar{I}} = 1$ as long as the marginal return to working capital investment $r_{\bar{I}}(w) \geq E[\tilde{r}_m]$ because working capital has a higher expected return and lower risk over this range than the risky asset, and because investing in inventory also dominates the risk-free asset since $E[\tilde{r}_m] > r_f$. In Section 5, we show that our calibrated model delivers sufficiently high marginal returns that this is the case at low levels of wealth. At higher levels of wealth where $r_{\bar{I}}(w) < E[\tilde{r}_m]$ the optimal allocation depends on the utility function, but as long as $r_{\bar{I}}(w) > r_f$ consumers will optimally split assets between working capital and the risky investment, as the risk-free bond is strictly dominated by the (approximately) risk-free investment in household inventory. As wealth becomes large, consumers will allocate all additional wealth to financial assets. Consequently, $\lambda_{\bar{I}}$ gradually declines as wealth increases.

5. Financial net returns to household inventory investment

Solving the optimization problem (23) yields the average monthly cost $V(\bar{I})$ of supplying consumption flow $\sum_l C_l$. To compute the marginal return to household working capital, we compute this cost at each level of household working capital $\bar{I}$.

In principle, we can then compute the marginal return as $-V'(\bar{I})$, providing a net return measure which incorporates both the price paid in store and also trip fixed costs and depreciation costs. In practice, the cost function is not smooth because $s_{l,d}$ is discrete. Consequently, when computing the marginal return in (24) numerically, $-\frac{V(\bar{I}+\Delta\bar{I})-V(\bar{I})}{\Delta\bar{I}}$ may be zero when $\Delta\bar{I}$ is small, but substantial when the increment is increased. We therefore utilize somewhat larger working capital increments of $\Delta\bar{I} = 2.5\%$ of annual consumption. This choice of $\Delta\bar{I}$ allows us to accurately characterize the marginal return while avoiding non-monotonicities (which we do not think are useful for thinking about real world behavior). Our results are robust to choosing somewhat larger or smaller increments.

Similarly, because the value function (23) is not well defined at zero working capital and positive trip fixed costs, we define the average (net) return relative to a low but non-zero working capital benchmark of $\bar{I}_0 = 2.5\%$ of annual consumption:

$$\text{Average (net) return:} \quad \bar{r}_{\bar{I}}(\bar{I}) = -\frac{V(\bar{I}) - V(\bar{I}_0)}{\bar{I} - \bar{I}_0}. \tag{30}$$

We choose not to report marginal returns for lower values of working capital because returns are extreme and unlikely to be relevant for most households (note that adjusting $\bar{I}_0$ does not have any effect on the marginal return at higher levels of working capital). At our choice of $\bar{I}_0$, the household only has just enough working capital to facilitate a weekly shopping trip with no stockpiling. Below this point, returns largely reflect households being able to shop at a lower frequency.

5.1. Model results of calculating financial net returns

Table 4 shows how increasing the maximum household working capital $\bar{I}$ affects the different sources of household savings, optimal trip interval Δ^* and stockpiling strategy $s^*_{l,d}$ (for the most durable products in group 4), and marginal and average financial net returns. In equilibrium, households do not stockpile products in groups 1 and 2 because the calibrated shelf-life is too short relative to the average trip interval.

When the amount of funds allocated to household working capital is low, the household is restricted in its ability to take advantage of deals and must choose a low value for $s_{l,d}$. This is because stockpiling products well in advance of when they are needed (i.e., a large $s_{l,d}$) is working capital intensive. As the working capital investment increases, households choose progressively higher values of $s_{l,d}$ for storable products. At the same time, an increase in working capital also allows households to spend more per trip, increasing the trip interval, reducing trip fixed costs, and raising bulk savings. This effect is strongest at very low levels of working capital. Given that we match the average NCP trip interval of about one week, a relatively small amount of working capital is required to achieve the desired trip interval and the value of working capital allocated to cash is fairly small.[17]

The trip interval and savings of each type need not be monotonic in $\bar{I}$. Relaxing the constraint may have positive or negative effects on these variables depending on whether the deal-focused or bulk-focused strategies dominate. If the household chooses to use the additional funds to make larger trips, this makes it more costly to buy items several trips in advance and can therefore lead to a reduction in stockpiling and deal savings. Alternatively, if the household uses the additional funds to increase stockpiling, this can put downward pressure on trip size due to depreciation costs and reduce bulk savings.

At low levels of household working capital investment, the marginal return to additional investment is very high. When household working capital is equal to 5% of annual consumption, the marginal return is around 55%. The marginal return gradually diminishes and reaches zero when household working capital is around one third of annual consumption. As we discuss in Section 4.5, households in the model do not participate in the stock market if the marginal return to working capital is more than the expected stock market return. Over the two decades prior to our sample, the average annual S&P 500 return was around 8%. In Table 4, this corresponds to inventory share cutoffs of between 10.7% and 13.6% of annual spending.

When aggregating to the product group level, around 7% of households have an inventory ratio below this cutoff. The share increases to 71% when aggregating to department. In addition to uncertainty due to aggregation assumptions, we also expect this participation cutoff to vary substantially across households in our sample (for example, it is influenced by household-specific investment opportunities, beliefs about future investment returns, shopping trip fixed costs, and preferences for consuming perishable goods). It is therefore difficult to infer the share of households for whom our model predicts non-participation based purely on our inventory measure. Furthermore, some share of inventory holdings is explained by non-financial factors such as goods stockpiled for emergencies. To the extent that these other factors raise the value of inventories, more than 7% of households could be below the non-participation cutoff for the inventory ratio.

5.2. For which households is working capital important?

In this section we relate our model conclusions to observable household characteristics in the NielsenIQ data. There are two distinct ways to think about the importance of working capital as an asset. Firstly, working capital may be considered to be important for a household if it can rationalize stock market non-participation. This suggests we should look for the characteristics that predict a high *marginal* return to working capital (i.e., low working capital relative to consumption). However, an arguably more important question is how working capital affects overall portfolio returns. For portfolio returns, what matters is the *average* return to working capital and the ratio of working capital to financial assets. A low *marginal* return is quite consistent with a large effect on portfolio returns. We examine each of these questions in turn.

5.2.1. Who has low inventory in practice (and do they have high marginal returns)?

Understanding which households earn high marginal returns is relevant for a discussion of the stock market participation puzzle as these are the households for whom we may be able to rationalize non-participation. In Section 5.1, we noted that the model implied an inventory-to-spending ratio cutoff of 10.7–13.6% for stock-market non-participation. Unfortunately, it is difficult to directly map the participation cutoff to inventory-to-spending ratios we observe in the data. Instead we repeat the exercise from Section 3.4, predicting whether households have an inventory ratio in the bottom quartile or not. Figure H.5 shows the relative importance of each predictor. The top six characteristics are the same as when we predict the inventory ratio: the expenditure share of perishable goods, store type shares (which may also capture perishability), age, density, and house prices. Figure H.6 shows partial dependence profiles.

Linking this to the model, we expect older households have a lower value of k (low cost of time) and the perishable share enters directly (C_1). It is plausible that the grocery store share and other store type shares are also linked to product storability. Households with a high share of grocery spending are much more likely to have low inventory. Lower k and higher C_1 reduce the working capital ratio cutoff associated with stock market participation, all else equal.

5.2.2. What is the effect of working capital on portfolio return heterogeneity?

There do not appear to be large differences by income in the degree to which households exploit returns to working capital. However, working capital is much more important for low income and low wealth households because it is a larger share of their *overall portfolio* and has higher average returns than traditional financial assets. Including working capital alongside financial assets may dramatically change estimates of portfolio return heterogeneity. To quantify this, we use the SCF to compute annual portfolio shares and returns to each type of asset, assuming a household's capital gain for a given asset class is equal to that asset's aggregate capital gain.

We then assign a level of inventory to each SCF household based on a predictive model computed using the NielsenIQ data, approximating the working capital share using the inventory share. We assign to working capital an average return of 54 per cent, the average return in Table 4 corresponding to the average NielsenIQ inventory-to-spending ratio of around 20% in Column 2.

Fig. 9 plots the average portfolio return for households in each income and asset quartile, with and without working capital. As has been documented elsewhere (e.g., Fagereng et al. (2020) or Bach et al. (2020)), average returns on financial assets tend to increase with both income and assets. Because working capital is a large share of low income households' assets and has a high average return, including it changes this pattern dramatically. It seems likely that incorporating working capital would increase the average returns of low income households substantially relative to other households.[18]

5.3. An empirical measure of returns based on in-store savings data

Previously, we used the calibrated model to compute the *net* return to working capital (marginal and average) as a function of working

[17] The level of cash and cash-equivalent asset holdings predicted by the model should of course not match the level observed in a comprehensive household finance survey since the model only captures one motive for holding cash and leaves out other motives such as precautionary liquidity. Furthermore, our model applies to other goods not covered by the NielsenIQ data which require additional cash holdings.

[18] Appendix Figure H.7 shows that including retirement accounts reduces the effect of working capital on average returns at higher levels of income and assets, but does not change the overall conclusion.

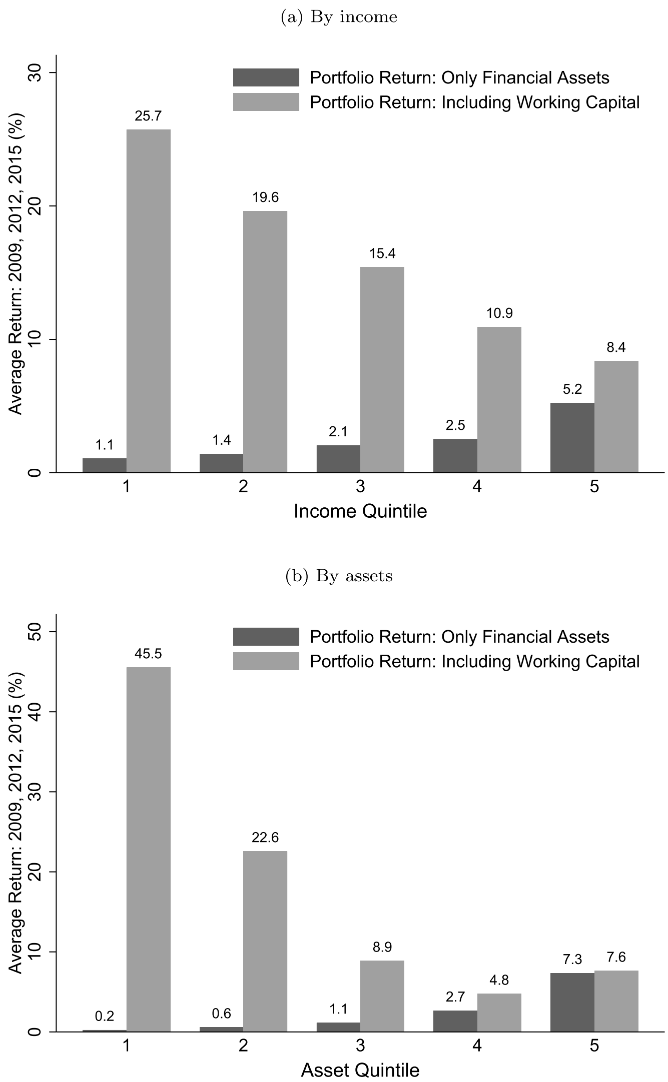

Fig. 9. Implied effect of working capital on average returns.
Notes: This graph shows average portfolio returns in each income and asset quintile with and without working capital. We construct average returns for each SCF household using the following method. First we compute the portfolio return without working capital, which is $r_i = \frac{\text{Income from Financial Assets}_i}{\text{Financial Assets}_i} + \sum_a \lambda_{a,i}\text{Revaluation}_a$, where $\lambda_{a,i}$ is household i's share of assets in asset class a and Revaluation_a is the revaluation return for asset class a from the Flow of funds (the SCF only provides information on realized capital gains and losses). Income from financial assets is pre-tax and includes both interest and dividend income. The assets classes are stock mutual funds, directly held stocks, bond mutual funds, directly held bonds and combined mutual funds. Also included in financial assets are checking accounts, savings accounts, CDs and money market accounts. We assume these assets have zero revaluation return. To compute the corresponding return including working capital we use the inventory portfolio share $\lambda_{I,i} = \frac{\text{Inventory}_i}{\text{Assets}_i + \text{Inventory}_i}$ where Inventory_i is imputed inventory of household i (see note to Fig. 2). The return including working capital is then $r_i^{wc} = \lambda_{I,i}\text{AvgWorkingCapitalReturn}_{q(i)} + (1 - \lambda_{I,i})r_i$, where $q(i)$ is the income quintile of household i. Note that as we observe inventory (not working capital), we ideally need the average return to inventory (not working capital). However, at the average inventory ratio of 0.2 these are essentially identical. We therefore use the average return to working capital associated with an inventory-to-spending ratio of around 20%, which is 54% (Table 4). Income is reported to the nearest thousand dollars. The lower cutoffs for each income quintile are \$0, \$22,000, \$38,000, \$61,000 and \$101,000. The lower cutoffs for each asset quintile are \$0, \$321, \$2,001, \$8,510, \$43,600. All quintiles and summary statistics are calculated using SCF weights.

capital, which we cannot recover directly from the data. These return functions were obtained for a single model household that represents the typical household in the data. An alternative measure of returns which, in principle, we can compute using the data alone is the 'gross return', which ignores trip fixed costs and holding or depreciation costs. This measure of returns reflects in-store savings only, and we can in principle estimate it using variation in working capital across households in the data. While we do not observe working capital, this is only a minor limitation because unobserved cash held to facilitate shopping trips makes up only a small share of model working capital at inventory-to-spending ratios observed in the data; see column 3 of Table 4 and Fig. 1. We therefore use inventory to measure working capital $\bar{I}_h$.[19]

One issue with estimating the relationship between in-store savings and working capital directly is that in the data—contrary to the model—all else is not held constant. The most obvious problem is that households with a higher level of overall spending have both a higher dollar amount of inventory and a higher dollar amount of savings. A natural solution is to divide both annual average dollar in-store savings and annual average inventory by annual average spending before estimating the average gross return as the parameter b in this relationship[20]:

$$\frac{\text{Dollar in-store savings}_h}{\text{Spending}_h} = a + b\frac{\text{Average Inventory}_h}{\text{Spending}_h} + e_h. \quad (31)$$

However, holding consumption fixed, households who obtain lower prices and thus higher in-store savings also have mechanically lower spending and a lower level of inventory because their average purchase price is lower. To address these potential sources of bias, we compute alternative measures of spending and inventory at fixed prices, which we refer to as "base spending", "base inventory", and "base price". Our definition of base price corresponds to the expected price in the model when $s_{l,d}=0$ and $b(Q_l)=1$. That is, the average price paid for a product if the household engages in "untargeted shopping" in their location and buys the "standard" pack size. We define base spending as the amount the household would have spent if they had purchased an identical basket of items at a fixed base price. Base inventory is also computed using a fixed-price measure of spending. The construction of base spending is discussed in more detail in Appendix A and base inventory in Appendix E.2.

A related issue is that trip fixed costs are held constant for the model household, but are obviously not constant across households in the data. Since they are an important driver of inventory in the model, unobserved differences in trip fixed costs k can lead to differences in inventory even when holding working capital $\bar{I}$ fixed. For instance, high trip fixed costs reduce deal discounts in the model, generating a negative relationship between inventory and deal savings when varying k rather than $\bar{I}$. To control for the effect of unobserved trip fixed costs, we therefore condition on the household's average trip interval when estimating the relationship between inventory and savings. We discuss the data relationships between inventory, the trip interval, and savings in more detail in Section 5.4.[21]

[19] Note that while cash and inventory are negatively correlated over the trip cycle in the model, we estimate the relationship between *annual average* inventory and *annual average* in-store savings.

[20] In interpreting b as the gross return we are assuming that in-store savings increase linearly with working capital. If true, this in turn implies that the marginal gross return is constant and equal to the average gross return. While this is not the case in our model, it is a reasonable approximation in the data over the range of inventory ratios we observe.

[21] Controlling for the trip interval removes most of the variation in bulk savings in the model, and some of the variation in the data. Our gross returns measure may therefore be more accurately interpreted as capturing returns through stockpiling rather than by extending the trip interval and buying in bulk. In the model, controlling for the trip interval is not necessary for identification because the trip fixed cost is held fixed, but it is necessary for a consistent

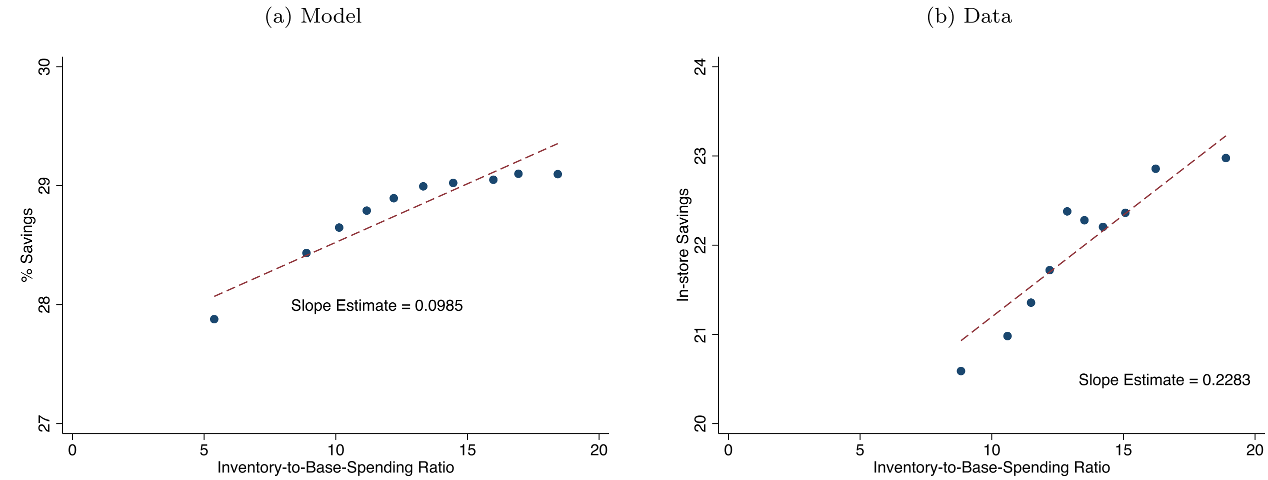

Fig. 10. Gross return: in-store savings and inventory.
Notes: To construct Panel (a) we vary working capital ($\bar{I}$) and compute the in-store savings and inventory as a share of spending. We plot the relationship between the two conditional on the number of trips for comparability with the data. Panel (b) uses the NCP over 2013 and 2014 to illustrate the relationship between in-store savings and average inventory as a percentage of spending. The points on the charts represent deciles of households, controlling for the number of shopping trips a household makes each year. The savings measure reflects in-store savings only and does not incorporate holding costs or trip fixed costs. The red dashed line shows predicted values from (32): $\frac{\text{Dollar in-store savings}_h}{\text{Base Spending}_h} = a + b_1 \frac{\text{Average Base Inventory}_h}{\text{Base Spending}_h} + b_2 \text{Trip Interval}_h + e_h$. The regression is weighted using NielsenIQ sampling weights.

Taking into account these considerations, we measure (conditional) gross returns as b_1 in the following regression specification:

$$\frac{\text{Dollar in-store savings}_h}{\text{Base Spending}_h} = a + b_1 \frac{\text{Average Base Inventory}_h}{\text{Base Spending}_h} + b_2 \text{Trip Interval}_h + e_h. \quad (32)$$

Fig. 10 shows estimates of b_1 from (32) and a binned scatterplot for both the model and the data.[22] Holding the trip interval fixed, varying working capital generates a positive relationship between savings and inventory in both the model and the data. The gross return observed in the model declines with the inventory ratio. The return over the range of inventory ratios we consider is around 10%. Our data return of 23% is higher than the model, although Appendix Figure H.9 shows that gross returns are lower when aggregating to product module and when aggregating to department.

5.3.1. *Moving events as an alternative source of variation to measure Gross returns*

One reason why model and estimated gross returns potentially differ is that the estimation of b_1 in (32) might use variation in inventory that is not exogenous.[23] A partial solution to this problem is found by leveraging the variation driven by households that move locations. Because it is costly to move a stockpile of household goods, we expect that households will reduce s_d in advance of a move, consistent with Fig. 3. Informally, we expect this will also lead to a reduction in deal savings obtained in the weeks prior to a move as households reduce stockpiling items on sale. In addition, as households who recently moved have a limited stockpile, we expect that a larger share of items will be purchased at full-price after the move as well.

Fig. 11 shows that the share of deal transactions indeed drops by 3.6 percentage points in the month of the move and recovers only gradually in the months following the move, consistent with this conjecture. We focus on the share of deal transactions rather than percentage savings, as our savings measure can only be computed when the exact NCP UPC-store-week combination corresponds to an observation in the NRP. Because the coverage varies substantially from month to month at the household level, the resulting monthly household savings is too noisy to credibly estimate effects around the move. In contrast, the NCP deal indicator is available for all NCP transactions. Appendix Figure H.10 shows that a one percentage point increase in the NCP deal indicator is associated with an 0.15% increase in percentage deal savings on a consistent set of transactions. The 3.6 percentage point drop in Fig. 11 therefore corresponds to an 0.54 percentage point (or about 10%) drop in deal savings.

Combining this with an implied change in the inventory ratio of 6.63% of annual spending (see Fig. 3) gives a back-of-the-envelope gross return of $100 \times \frac{0.54}{6.63} = 8.1\%$.[24] As this approach uses within household variation, we do not condition on the trip interval and therefore compare our data gross return to the unconditional model return of 12.2% shown in Appendix Figure H.8. The movers estimate is lower than the cross-sectional estimates and fairly comparable to the gross return implied by the model.

5.4. *Model validation and robustness*

In the model, households obtain high returns to working capital by exploiting temporary sales and bulk discounts. The model generates a number of predictions for the relationships between working capital $\bar{I}$,

comparison with the data relationship. Appendix Figure H.8 shows that the unconditional model gross returns are in any case mostly driven by deal savings, except at very low levels of inventory. Bulk savings are important at low levels of inventory because households who are constrained to small trip intervals and pack sizes pay much higher prices.

[22] Dividing by base spending reduces the average inventory ratio relative to what we report in Section 3 because base spending is higher than observed spending.

[23] Alternative sources of variation in inventory which are present in the model are trip fixed costs k, which we can control for by conditioning on the trip interval, and variation in holding costs. In the data, variation in holding costs may reflect differences in the space available to store items, preferences for perishable products, or in the way grocery items are processed and stored. In the model, variation in δ_l generates a similar relationship between inventory and deal savings as variation in $\bar{I}$. Variation in holding costs across households may therefore also contribute to the data relationship in Fig. 12d. In addition, there are likely sources of inventory variation in the data that are not present in the model.

[24] There is not a statistically significant change in bulk savings around a move so we assume the response is zero. This is consistent with bulk savings driven by a high level of consumption of a particular product.

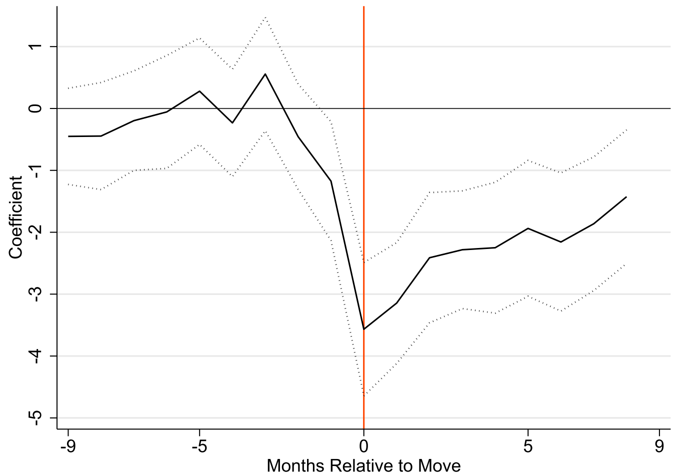

Fig. 11. Deal transactions around move dates.
Notes: This figure shows the percentage point change in the share of purchases associated with a coupon or self-reported deal. For households who move to a new 3-digit ZIP Code in a given year we identify the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather than their old 3-digit ZIP Code). The figure plots estimates of b_s from the following specification and a 95 per cent confidence interval: $\text{Deal Share}_{h,t} = \sum_{s=-9}^{9} b_s \text{ Moved}_{h,t-s} + \text{Month FE} + \text{Household FE} + e_{h,t}$, The sample includes both non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. Standard errors are clustered by household. The regression is weighted using NielsenIQ sampling weights. Appendix Figure H.11 shows robustness to using the imputation approach described by Borusyak et al. (2021) to deal with potential bias in pooled event studies with staggered events.

the trip fixed cost k, and in-store savings, which we can test using the NCP, in addition to the tests we performed in the previous Section 5.3 and in Section 3.3. These tests help support our claim that financial returns to working capital are an important determinant of households' high inventory holdings.

The first exercise we perform using the model is to show the effect of the fixed cost k on savings. That is, we vary k and plot the resulting relationship between in-store savings and the trip interval Δ. In the second exercise, we vary working capital $\bar{I}$ and plot the relationship between savings and average inventory generated by this variation. To compare these relationships with the data, we construct measures of deal savings and bulk savings using the NCP and NRP using the base spending and base inventory measures from Section 5.3.

To look at the effect of k on savings, we estimate the following equation separately for deal savings and bulk savings, and for the NCP and data generated from the model by varying k:

$$\frac{\text{Dollar in-store savings}_h}{\text{Base Spending}_h} = a + b_1 \text{Trip Interval}_h + e_h. \tag{33}$$

To look at the effect of varying $\bar{I}$ on savings, we estimate (32) separately for deal savings and bulk savings, and for the NCP and data generated from the model by varying $\bar{I}$.[25]

5.4.1. *Deal savings*

We measure deal savings as the discount relative to the amount the household would have spent if they paid the average price for the same UPC in the same store over that year: the *additional* savings resulting from strategic shopping behavior relative to random shopping in the same store over time. Note that because different pack sizes of the same product have different UPCs, this measure does not incorporate bulk savings. We discuss the construction of the deal savings measure in more detail in Appendix A.

Fig. 12a plots the model relationship between deal savings and the trip interval generated by varying k. Deal savings are lower for households with longer trip intervals (higher fixed costs). Fig. 12b shows that deal savings are also lower for households with longer trip intervals in the NCP. This is consistent with Aguiar and Hurst (2007): households in the NCP who shop more frequently obtain higher in-store savings. In our model, shopping more frequently allows households to set a higher value of s_d, holding working capital fixed, and therefore obtain higher deal savings because prices are observed more frequently.

Next, we vary working capital $\bar{I}$ and plot the model relationship between in-store savings and the ratio of inventory to annual spending. Fig. 12c shows that increasing working capital allows households in the model to obtain more savings by allowing households to set a higher value of s_d. Fig. 12d shows that we also observe a positive relationship between inventory and deal savings in the data. Allocating more working capital to use for stockpiling can be seen as a substitute for shopping more frequently. Consistent with this, returns to working capital in the model are higher for households with a high cost of time (i.e., a high k). Households with a low cost of time, such as the older households studied by Aguiar and Hurst (2007) are able to exploit most of the potential deal savings with a small level of working capital because they shop very frequently.

5.4.2. *Bulk savings*

Our measure of bulk savings compares the average unit price paid with the average unit price of a second quintile pack size of the same product in the same year and 3-digit ZIP Code. We describe our method for computing bulk savings in more detail in Appendix A. We add an additional control when estimating the data relationships between bulk savings, the trip interval and inventory to improve comparability with the model. In practice, the potential for bulk savings varies across products—some products are available (and much cheaper) in larger pack sizes, whereas others are not. We therefore define potential bulk savings as the bulk savings that would be obtained if the household bought every product in the largest pack size available in their 3-digit ZIP Code. Potential bulk savings explains around 65% of the variation in observed bulk savings.

Fig. 13a plots the model relationship between trip interval and bulk savings. In the model, households with high fixed costs k, and therefore longer trip intervals Δ, obtain more bulk savings because households with less frequent trips also buy more each trip. While the data relationship is also positive, the slope is close to zero (Fig. 13b).

There is also a flat relationship between the inventory ratio and bulk savings in the model (Fig. 13c). This follows from our assumption that bulk savings are obtained by choosing a larger steady state trip size. Consequently, variation in bulk savings generated by adjusting $\bar{I}$ comes purely through the effect on the trip interval, which we controlled for in this specification for the reasons outlined in Section 5.4.1. Fig. 13d shows that the relationship in the data is similar, supporting our assumption. That is, in both the data and the model, stockpiling mostly reflects households taking advantage of temporary deals on standard pack sizes, rather than buying in bulk. The fact that we find no change in bulk savings around a move is also consistent with this interpretation. One possible reason for this is that depreciation is often much higher after the pack is opened; buying many small packs will reduce depreciation costs.

[25] Note that we do not condition on the inventory ratio when estimating the relationship between the trip interval and savings. This is because when holding $\bar{I}$ fixed, conditioning on the inventory ratio absorbs essentially all the variation in the trip interval generated by k in the model.

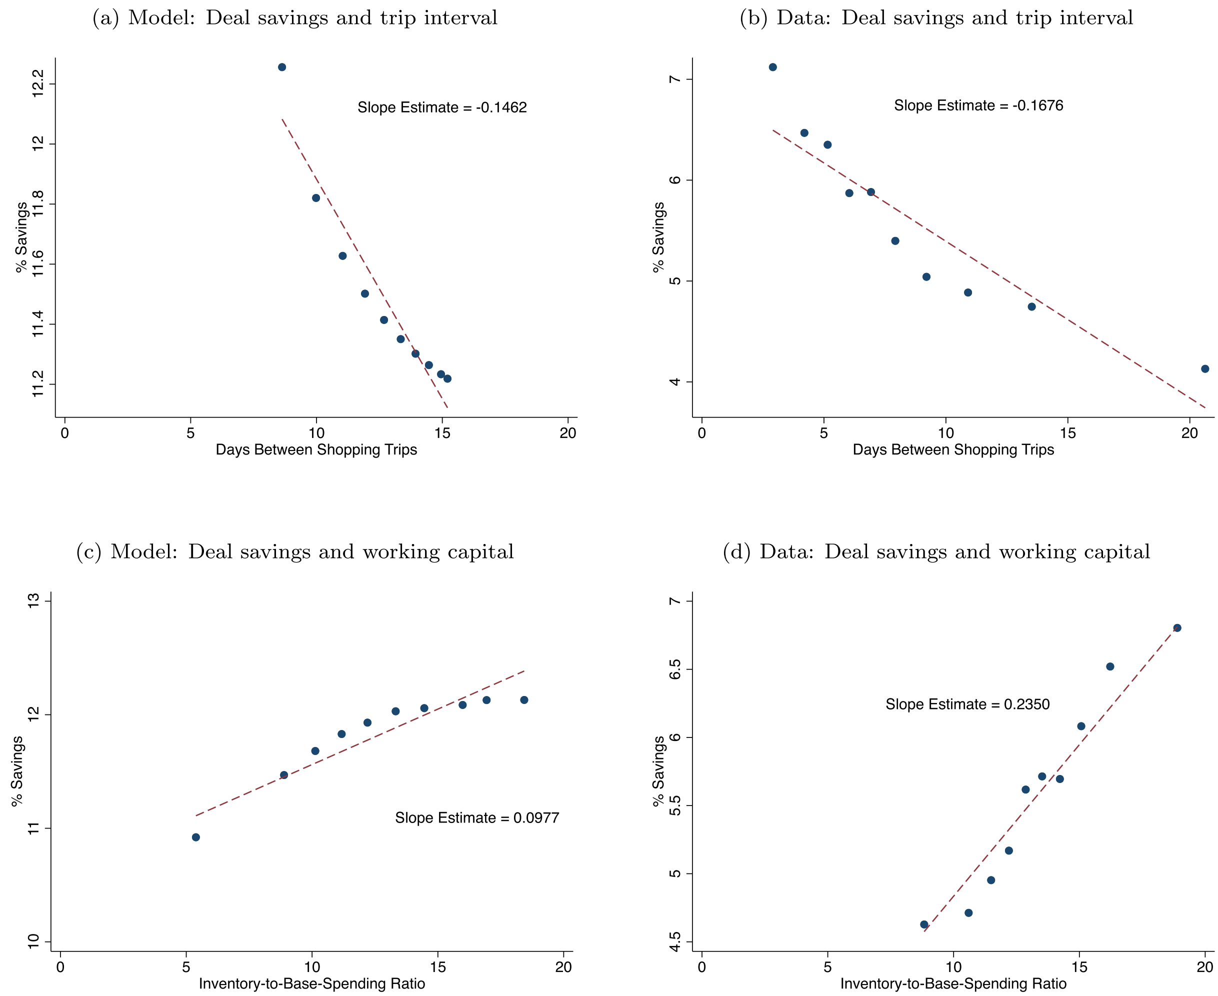

Fig. 12. Relationship between deal savings, trip interval, and working capital.
Notes: In Panel (a) we evaluate % deal savings and days between trips in the model for different values of k. Working capital is set sufficiently high that the working capital constraint does not bind. Days between trips are computed as $\Delta \times \frac{365}{12}$. In Panel (b) we plot the data relationship between deal savings and the household's average number of days between trips over a calendar year, winsorized at 98 per cent. In Panel (c) we evaluate % deal savings and the inventory ratio in the model for different values of $\bar{I}$. Panel (d) shows the corresponding relationship in the data. The inventory ratio in the model is average inventory divided by annual spending at the expected price when $s_d = 0$ and $b(Q) = 1$. The inventory ratio in the data is the household's average base inventory over a calendar year divided by annual NielsenIQ base spending (see Appendix A and E.2 for an explanation of how these variables are constructed). Deal savings in the data are constructed using equation 44 in Appendix A. In Panels (c) and (d) we control for the number of trips. Panels (b) and (d) are constructed using NielsenIQ sampling weights.

6. Conclusion

We study how households can obtain substantial financial returns from strategic shopping behavior and optimally managing inventories of consumer goods. We find that American households tend to hold substantial amounts of these non-financial assets and rationally choose to maintain some amount of liquid savings not only for precautionary motives but in support of this inventory management role. Such inventories are missing from traditional consumer finance data such as the SCF, which might explain why household working capital has been largely ignored by the household finance literature.

We demonstrate that households earn high marginal returns from inventory management through several channels at low levels of inventory, but these marginal returns decline rapidly as inventory levels increase. At low levels of inventory, the marginal return to investment in inventory strongly dominates stock market returns and then quickly approach zero. Hence, even though marginal returns to working capital investment are low for households that are not borrowing constrained, average returns are still high and contribute substantially to average portfolio returns, especially for low income and low wealth households.

Though we do not consider them explicitly in the paper, time-varying investment opportunities in working capital such as large temporary store price discounts, sales tax holidays, or 'Black Friday' sales could even rationalize some of the borrowing at high interest rates that lower income households engage in (Zinman 2015). Investment in working capital is therefore related to the literature motivating household borrowing as a way to invest in illiquid assets offering high rates of return but requiring a threshold amount of capital (Angeletos et al. 2001; Laibson et al. 2003).

Since we do not observe financial assets and borrowing costs and limits of households in the scanner data, our model of optimal household inventory management does not feature credit constraints. We therefore view the collection of more comprehensive household balance sheet data, including household working capital and borrowing limits, an important next step in this line of research.

Finally, we note that adding household inventory management to a household's portfolio choice problem can potentially affect its decision of whether to participate in financial markets. Household working capital could therefore provide another partial explanation to the

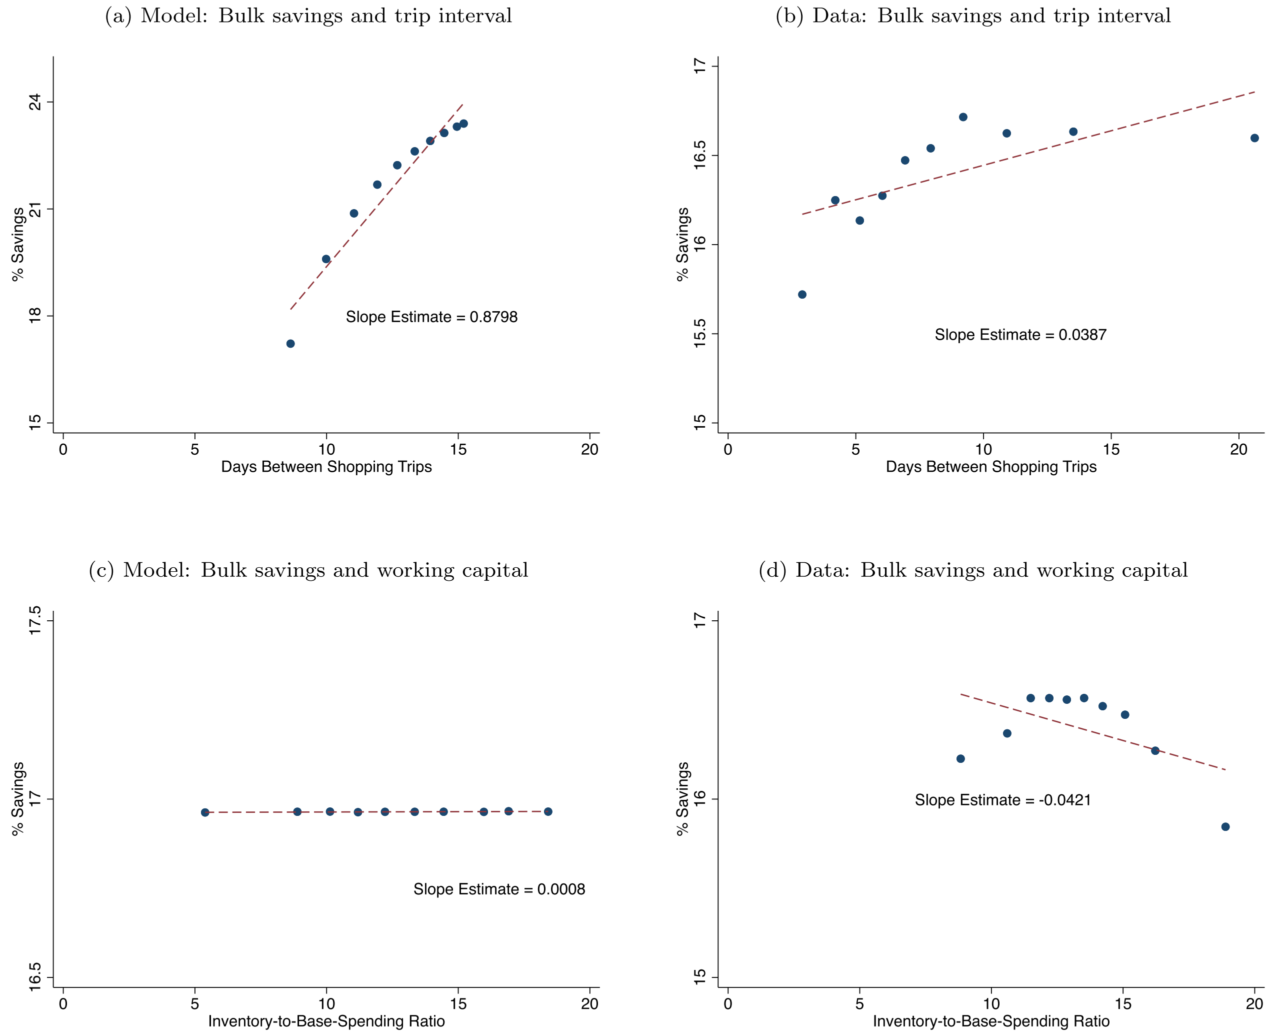

Fig. 13. Relationship between bulk savings, trip interval, and working capital.
Notes: In Panel (a) we evaluate % bulk savings and days between trips in the model for different values of k. Working capital is set sufficiently high that the constraint does not bind. Days between trips are computed as $\Delta \times \frac{365}{12}$. In Panel (b) we plot the data relationship between bulk savings and the household's average number of days between trips over a calendar year, winsorized at 98 per cent. In Panel (c) we evaluate % bulk savings and the inventory ratio in the model for different values of $\bar{I}$. Panel (d) shows the corresponding relationship in the data. The inventory ratio in the model is average inventory divided by annual spending at the expected price when $s_d = 0$ and $b(Q) = 1$. The inventory ratio in the data is the household's average base inventory over a calendar year divided by annual NielsenIQ base spending (see Appendix A and E.2 for an explanation of how these variables are constructed). Bulk savings in the data are constructed using equation 46 in Appendix A. In Panels (c) and (d) we control for the number of trips. In Panels (b) and (d) we also control for potential bulk savings, which is the bulk savings obtained if the household purchased the largest available pack size quintile for each product. Panels (b) and (d) are constructed using NielsenIQ sampling weights.

stock market non-participation puzzle.[26] Since investment in household working capital has investor-specific and approximately risk-free returns that decline systematically as wealth increases and that dominate equity returns for poorer households, it complements explanations of cross-sectional variation in participation rates, such as participation costs. However, at this point these are conjectures and open to future research as we do not observe financial asset holdings in the scanner data.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code and data packet for this article can be found at https://doi.org/10.17632/66w4y6sb6d.1.

[26] The literature on the participation puzzle is among the oldest in household finance and too large to adequately survey here. The handbook chapters by Guiso and Sodini (2013) and Beshears et al. (2018) provide recent surveys. A non-exhaustive list of explanations of the participation puzzle include pecuniary and non-pecuniary participation fixed costs (Luttmer 1999; Vissing-Jørgensen 2002); low financial literacy (Van Rooij et al. 2011; Black et al. 2018); non-expected utility with first-order risk aversion (Barberis et al. 2006; Epstein and Schneider 2010); heterogeneity in beliefs (Kézdi and Willis 2009; Malmendier and Nagel 2011; Hurd et al. 2011; Adelino et al. 2020), lack of trust (Guiso et al. 2008; Gennaioli et al. 2015); unawareness of the excess return premium (Guiso and Jappelli 2005; Grinblatt et al. 2011; Cole et al. 2014); background risk (Heaton and Lucas 2000; Cocco et al. 2005) and positive correlation of stock returns with returns of other assets in portfolios (Benzoni et al. 2007; Davis and Willen 2014; Bonaparte et al. 2014); liquidity constraints, illiquid assets and consumption commitments (Grossman and Laroque 1990; Haliassos and Michaelides 2003; Chetty and Szeidl 2007); or social interactions (Hong et al. 2004; Kaustia and Knüpfer 2012).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jfineco.2023.103758.

References

Adelino, M., Schoar, A., Severino, F., 2020. Perception of house price risk and homeownership. NBER Working Paper No. 25090.

Aguiar, M., Hurst, E., 2007. Life-cycle prices and production. Am. Econ. Rev. 97, 1533–1559.

Aguiar, M., Hurst, E., 2013. Deconstructing life cycle expenditure. J. Polit. Econ. 121, 437–492.

Angeletos, G.M., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S., 2001. The hyperbolic consumption model: calibration, simulation, and empirical evaluation. J. Econ. Perspect. 15, 47–68.

Arrow, K.J., Harris, T., Marschak, J., 1951. Optimal inventory policy. Econometrica, 250–272.

Bach, L., Calvet, L., Sodini, P., 2020. Rich pickings? Risk, return, and skill in household wealth. Am. Econ. Rev. 110, 2703–2747.

Baker, S.R., Johnson, S., Kueng, L., 2021. Shopping for lower sales tax rates. Am. Econ. J. Macroecon. 13, 209–250.

Baker, S.R., Kueng, L., McGranahan, L., Melzer, B.T., 2019. Do household finances constrain unconventional fiscal policy? Tax Policy Econ.

Barberis, N., Huang, M., Thaler, R.H., 2006. Individual preferences, monetary gambles, and stock market participation: a case for narrow framing. Am. Econ. Rev. 96, 1069–1090.

Bartmann, D., Beckmann, M.J., 1992. Inventory Control: Models and Methods. Springer-Verlag.

Benzoni, L., Collin-Dufresne, P., Goldstein, R.S., 2007. Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. J. Finance 62, 2123–2167.

Bertaut, C.C., Starr, M., 2000. Household portfolios in the United States. FEDS Working Paper No. 2000-26.

Beshears, J., Choi, J.J., Laibson, D., Madrian, B.C., 2018. Behavioral household finance. In: Bernheim, D.B., DellaVigna, S., Laibson, D. (Eds.), Handbook of Behavioral Economics: Applications and Foundations, vol. 1. Elsevier, pp. 177–276.

Black, S.E., Devereux, P.J., Lundborg, P., Majlesi, K., 2018. Learning to take risks? The effect of education on risk-taking in financial markets. Rev. Finance 22, 951–975.

Boizot, C., Robin, J.M., Visser, M., 2001. The demand for food products: an analysis of interpurchase times and purchased quantities. Econ. J. 111, 391–419.

Bonaparte, Y., Korniotis, G.M., Kumar, A., 2014. Income hedging and portfolio decisions. J. Financ. Econ. 113, 300–324.

Borusyak, K., Jaravel, X., Spiess, J., 2021. Revisiting event study designs: Robust and efficient estimation. Working Paper.

Broda, C., Weinstein, D.E., 2010. Product creation and destruction: evidence and price implications. Am. Econ. Rev. 100, 691–723.

Browning, M., Crossley, T., 2001. The life-cycle model of consumption and saving. J. Econ. Perspect. 15, 3–22.

Browning, M., Crossley, T., 2009. Shocks, stocks, and socks: smoothing consumption over a temporary income loss. J. Eur. Econ. Assoc. 15, 1169–1192.

Chetty, R., Szeidl, A., 2007. Consumption commitments and risk preferences. Q. J. Econ. 122, 831–877.

Chevalier, J.A., Kashyap, A.K., Rossi, P.E., 2003. Why don't prices rise during periods of peak demand? Evidence from scanner data. Am. Econ. Rev. 93, 15–37.

Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. Rev. Financ. Stud. 18, 491–533.

Coibion, O., Gorodnichenko, Y., Hong, G.H., 2015. The cyclicality of sales, regular and effective prices: business cycle and policy implications. Am. Econ. Rev. 105, 993–1029.

Cole, S., Paulson, A., Shastry, G.K., 2014. Smart money? The effect of education on financial outcomes. Rev. Financ. Stud. 27, 2022–2051.

Davis, S.J., Willen, P., 2014. Occupation-level income shocks and asset returns: their covariance and implications for portfolio choice. Q. J. Finance 3, 1–53.

Einav, L., Leibtag, E., Nevo, A., 2010. Recording discrepancies in Nielsen homescan data: are they present and do they matter? Quant. Mark. Econ. 8, 207–239.

Epstein, L.G., Schneider, M., 2010. Ambiguity and asset markets. Annu. Rev. Financ. Econ. 2, 315–346.

Fagereng, A., Guiso, L., Malacrino, D., Pistaferri, L., 2020. Heterogeneity and persistence in returns to wealth. Econometrica 88, 115–170.

Fisman, R., Love, I., 2003. Trade credit, financial intermediary development, and industry growth. J. Finance 58, 353–374.

Food Safety and Inspection Service FoodKeeper Data, 2020. United States Department of Agriculture. Available at https://catalog.data.gov/dataset/fsis-foodkeeper-data. (Accessed 7 July 2020).

Gennaioli, N., Shleifer, A., Vishny, R., 2015. Money doctors. J. Finance 70, 91–114.

Griffith, R., Leibtag, E., Leicester, A., Nevo, A., 2009. Consumer shopping behavior: how much do consumers save? J. Econ. Perspect. 23, 99–120.

Grinblatt, M., Keloharju, M., Linnainmaa, J., 2011. Iq and stock market participation. J. Finance 66, 2121–2164.

Grossman, S., Laroque, G., 1990. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. Econometrica 58, 25–51.

Guiso, L., Jappelli, T., 2005. Awareness and stock market participation. Rev. Finance 9, 537–567.

Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. J. Finance 63, 2557–2600.

Guiso, L., Sodini, P., 2013. Household finance: an emerging field. In: Constantinides, G.M., Harris, M., Stulz, R.M. (Eds.), Handbook of the Economics of Finance, vol. 2. Elsevier, pp. 1397–1532.

Haliassos, M., Michaelides, A., 2003. Portfolio choice and liquidity constraints. Int. Econ. Rev. 44, 143–177.

Hall, G., Rust, J., 2007. The (s,s) policy is an optimal trading strategy in a class of commodity price speculation problems. Econ. Theory 30, 515–538.

Heaton, J., Lucas, D., 2000. Portfolio choice and asset prices: the importance of entrepreneurial risk. J. Finance 55, 1163–1198.

Hendel, I., Nevo, A., 2006a. Measuring the implications of sales and consumer inventory behavior. Econometrica 74, 1637–1673.

Hendel, I., Nevo, A., 2006b. Sales and consumer inventory. Rand J. Econ. 37, 543–561.

Hendel, I., Nevo, A., 2013. Intertemporal price discrimination in storable goods markets. Am. Econ. Rev. 103, 2722–2751.

Hong, H., Kubik, J.D., Stein, J.C., 2004. Social interaction and stock-market participation. J. Finance 59, 137–163.

Hurd, M., Van Rooij, M., Winter, J., 2011. Stock market expectations of Dutch households. J. Appl. Econom. 26, 416–436.

Jappelli, T., Guiso, L., Haliassos, M., 2002. Household Portfolios. MIT Press.

Kaplan, G., Menzio, G., 2016. Shopping externalities and self-fulfilling unemployment fluctuations. J. Polit. Econ. 124, 771–825.

Kaplan, G., Schulhofer-Wohl, S., 2017. Inflation at the household level. J. Monet. Econ. 91, 19–38.

Kaustia, M., Knüpfer, S., 2012. Peer performance and stock market entry. J. Financ. Econ. 104, 321–338.

Kézdi, G., Willis, R.J., 2009. Stock market expectations and portfolio choice of American households. University of Michigan Working Paper.

Laibson, D., Repetto, A., Tobacman, J., 2003. A debt puzzle. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps. Princeton University Press, pp. 228–266. Chapter 11.

Luttmer, E.G., 1999. What level of fixed costs can reconcile consumption and stock returns? J. Polit. Econ. 107, 969–997.

Malmendier, U., Nagel, S., 2011. Depression babies: do macroeconomic experiences affect risk-taking? Q. J. Econ. 126, 373–416.

National Health and Nutrition Examination Survey, 2013–2014. Centers for Disease Control and Prevention. Available at https://www.cdc.gov/nchs/nhanes/index.htm. (Accessed 17 September 2020).

Nevo, A., Wong, A., 2019. The elasticity of substitution between time and market goods: evidence from the great recession. Int. Econ. Rev. 60, 25–51.

Nielsen Company Consumer Panel, 2013–2014. Kilts Center for Marketing at the University of Chicago Booth School of Business. Available at https://www.chicagobooth.edu/research/kilts/datasets/nielsen. (Accessed 2 December 2015).

Nielsen Company Retail MSR Scanner Data, 2013–2014. Kilts Center for Marketing at the University of Chicago Booth School of Business. Available at https://www.chicagobooth.edu/research/kilts/datasets/nielsen. (Accessed 2 December 2015).

Orhun, A.Y., Palazzolo, M., 2019. Frugality is hard to afford. J. Mark. Res. 56, 1–17.

Parker, J., 1999. The reaction of household consumption to predictable changes in payroll tax rates. Am. Econ. Rev. 89, 959–973.

Petersen, M.A., Rajan, R.G., 1997. Trade credit: theories and evidence. Rev. Financ. Stud. 10, 661–691.

Rampini, A., 2019. Financing durable assets. Am. Econ. Rev. 109, 664–701.

Samphantharak, K., Townsend, R.M., 2010. Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and. Corporate Financial Accounting. Cambridge University Press.

Sethi, S., Cheng, F., 1997. Optimality of (s, s) policies in inventory models with Markovian demand. Oper. Res. 45, 931–939.

Stroebel, J., Vavra, J., 2019. House prices, local demand, and retail prices. J. Polit. Econ. 127, 1391–1436.

Survey of Consumer Finances, 2010, 2013, 2016. Board of Governors of the Federal Reserve System. Available at https://www.federalreserve.gov/econres/scfindex.htm. (Accessed 20 September 2017).

Van Rooij, M., Lusardi, A., Alessie, R., 2011. Financial literacy and stock market participation. J. Financ. Econ. 101, 449–472.

Vissing-Jørgensen, A., 2002. Limited asset market participation and the elasticity of intertemporal substitution. J. Polit. Econ. 110, 825–853.

Yang, S.A., Birge, J.R., 2018. Trade credit, risk sharing, and inventory financing portfolios. Manag. Sci. 64, 3667–3689.

Zinman, J., 2015. Household debt: facts, puzzles, theories, and policies. Annu. Rev. Econ. 7, 251–276.