

NBER WORKING PAPER SERIES

FINANCIAL RETURNS TO HOUSEHOLD INVENTORY MANAGEMENT

Scott R. Baker
Stephanie G. Johnson
Lorenz Kueng

Working Paper 27740
<http://www.nber.org/papers/w27740>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2020, Revised April 2021

We thank Alex Butler, Hector Calvo-Pardo, James Choi, Giuseppe Di Giacomo, John Gathergood, Gustavo Grullon, Luigi Guiso, Michael Haliassos, Xavier Jaravel, Pierre Mabile, Soren Leth-Petersen, Jonathan Parker, Luigi Pistaferri, Tarun Ramadorai, Antoinette Schoar, and Ginyuan Yang for valuable feedback and our discussant Francesco d'Acunto and participants at seminars at Einaudi Institute for Economics and Finance (EIEF), National University of Singapore (NUS), Northwestern University, Rice University, University of Lugano (USI), CEPR European Conference on Household Finance 2020, Copenhagen New Consumption Data Conference 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 Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen or the National Bureau of Economic Research. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Scott R. Baker, Stephanie G. Johnson, and Lorenz Kueng. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Financial Returns to Household Inventory Management
Scott R. Baker, Stephanie G. Johnson, and Lorenz Kueng
NBER Working Paper No. 27740
August 2020, Revised April 2021
JEL No. D11,D12,D13,D14,E21,G11,G51

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 \$1,100 on average. Households can obtain significant financial returns from shopping strategically and optimally 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 with low levels of inventory earn high returns from investing in household working capital, well above 20%, though returns decline rapidly as inventory levels increase. We provide evidence from scanner and survey data that supports this conclusion. High returns from inventory management that are declining in wealth offer a new rationale for poorer households not to participate in risky financial markets, while wealthier households invest in both financial assets and household working capital.

Scott R. Baker
Kellogg School of Management
Northwestern University
2211 Campus Drive
Evanston, IL 60208
and NBER
s-baker@kellogg.northwestern.edu

Lorenz Kueng
Northwestern University
Kellogg School of Management
Suite 4394
2211 Campus Drive
Evanston, IL 60208
and NBER
lorenz.kueng@gmail.com

Stephanie G. Johnson
Jones Graduate School of Business
Rice University
6100 Main Street
Houston, TX 77005
stephgw Wallace@gmail.com

1 Introduction

While a large number of American households hold small amounts or even zero financial assets, all households hold at least some wealth in the form of consumer good 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. Aggregating across all Nielsen Homescan goods, we estimate that households hold approximately \$1,100 in consumer goods inventory at any given time, representing an unmeasured source of non-financial wealth. For households in the lowest quintile of household income, this inventory likely makes up a majority of total household wealth. Moreover, households can earn high returns through the maintenance of liquid savings and engaging in strategic shopping behavior.

In this paper, we study how the financial return to investment in inventories affects households' portfolio allocation and desire to hold liquid assets like cash and cash equivalent assets (such as checking accounts, transaction accounts, credit card lines of credit, etc.). We refer to these combined resources – the sum of cash and inventory – as *household working capital*. We show that for low levels of working capital, the marginal returns to inventory management are very high and dominate stock market returns. While returns are high at low levels of working capital, they decline rapidly with inventory holdings.

Optimal inventory management provides a rationale for households to hold sizable amounts of household working capital above and beyond the desire to maintain a buffer stock or precautionary source of savings. If low-asset households hold a large share of their assets in the form of inventory, these motives will be relevant for understanding the ability of such households to smooth consumption in response to temporary income shocks. The high returns observed in our data can, in some instances, rationalize high-cost borrowing like credit card debt.

Using scanner data from AC Nielsen and income and asset data from the Survey of Consumer Finances (SCF), we provide evidence in support of this new channel. In particular, we compute the total net returns to investment in household working capital. We go one step further than existing work (e.g. [Griffith, Leibtag, Leicester and Nevo \(2009\)](#); [Nevo and Wong \(2019\)](#)), 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. Moreover, we extend our framework to include the costs from product depreciation and the relation between the level of inventory holdings and differences in shopping trip fixed costs associated with different shopping behaviors.

We build a parsimonious model of inventory management to incorporate these additional components of returns to household working capital investments (product spoilage, trip costs, etc.) to compute individual net returns. These net returns are risk-free and well above 20% at low levels of working capital, though they decline rapidly as inventory levels increase.

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.

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 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 depreciation costs.

Existing models of deal shopping focus on individual products in a stochastic framework (e.g. [Boizot, Robin and Visser \(2001\)](#); [Hendel and Nevo \(2013\)](#)). In contrast, we focus on an aggregate deterministic steady state, where a constant fraction of goods are 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 independent across products, 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

By highlighting the role of household working capital for households’ portfolio allocation – especially for households with relatively 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 has received considerable attention in finance (e.g. [Petersen and Rajan \(1997\)](#); [Fisman and Love \(2003\)](#); [Yang and Birge \(2018\)](#)), inventories of consumer goods and household working capital has 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, Guiso and Haliassos \(2002\)](#) include household inventories. This also applies to the chapter by [Bertaut and Starr \(2000\)](#), who study U.S. households’ portfolios.³ One explanation for this gap is that inventories are often hard 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 estimate 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.⁴ We also highlight how adding household inventory management to a household’s portfolio choice problem affects its decision of whether to participate in financial markets. Household working capital therefore provides another partial explanation to the stock market par-

¹In 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. Stores generally feature consistent amounts of goods on sale throughout the year rather than concentrating deals in particular weeks; see Appendix Figures [B.1](#) and [B.2](#).

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

³The category “other non-financial assets,” which could in principle include inventories, does not. Instead, it includes “all standard passenger vehicles (cars, trucks, vans, minivans, jeeps, etc.) not owned by a business; all other types of personal-use vehicles (motor homes, recreational vehicles, planes, boats, motorcycles, etc.); and miscellaneous non-financial assets such as artwork, antiques, jewelry, furniture, and valuable collections (coins, stamps, etc.).”

⁴There is a related literature in macroeconomics studying heterogeneity in the effective price paid for similar goods across households and over business cycles; e.g. [Chevalier, Kashyap and Rossi \(2003\)](#); [Aguiar and Hurst \(2007\)](#); [Coibion, Gorodnichenko and Hong \(2015\)](#); [Kaplan and Menzio \(2016\)](#); [Kaplan and Schulhofer-Wohl \(2017\)](#); [Stroebel and Vavra \(2019\)](#).

ticipation puzzle: the fact that many households do not participate in risky financial assets to take advantage of the risk premium as predicted by standard portfolio theory (Mankiw and Zeldes (1991); Haliassos and Bertaut (1995)). The literature on the participation puzzle is among the oldest in household finance and too large to adequately survey here.⁵

We contribute to this literature by showing that investment in household working capital has returns that vary systematically across households by wealth and that dominate equity returns for poorer households. Importantly, the return on working capital is investor-specific, approximately risk-free, and declines as wealth and inventory holdings increase.

Hence, working capital investment with returns that decrease with wealth meets the challenge posed by Guiso and Sodini (2013) “to identify when and for which investors some of the explanations [of non-participation] are more relevant than others.” In this respect, this new explanation is comparable to participation costs as it applies to all households and is directly related to wealth. An interesting difference is that returns to working capital investment can be high because of (and not despite) the fixed costs involved, which in this case are the shopping trip fixed costs.

Even though we do not consider them explicitly in the paper, time-varying investment opportunities in working capital such as temporary large store price discounts, sales tax holidays or “Black Friday” sales could potentially rationalize the observation that poorer households often borrow at fairly high interest rates (e.g. Zinman (2015)). Investment in working capital is therefore related to the literature that motivates household borrowing as a way to invest in illiquid assets which offer high rates of return but require a small amount of capital to reach a certain threshold for investment, such as contributing to an employer-matched 401(k) retirement savings plan or making a down-payment for a home purchase (e.g. Angeletos, Laibson, Repetto, Tobacman and Weinberg (2001); Laibson, Repetto and Tobacman (2003)).

Alterations in strategic shopping behavior also help explain portions of the “excess sensitivity” of consumption to anticipated temporary income changes experienced by households (e.g. Jappelli and Pistaferri (2010), Kueng (2018)). In the existing consumption literature, many retail purchases, such as grocery store and pharmacy spending, are treated strictly as non-durables. As much of the literature moves to monthly and even daily measures of spending to improve identification of causal effects, the wedge between household spending and actual consumption grows more important. We show that households hold substantial stocks of consumer goods, make purchases in discrete bundles, and run them down over significant periods of time. This reinforces the idea that large increases in household spending in one period may translate into increased consumption only over

⁵The handbook chapters by Guiso and Sodini (2013) and Beshears, Choi, Laibson and Madrian (2018) provide recent surveys of this literature. 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, Lusardi and Alessie (2011); Black, Devereux, Lundborg and Majlesi (2018)); non-expected utility with first-order risk aversion (Barberis, Huang and Thaler (2006); Epstein and Schneider (2010)); heterogeneity in beliefs (Kézdi and Willis (2009); Malmendier and Nagel (2011); Hurd, Van Rooij and Winter (2011); Adelino, Schoar and Severino (2020)), lack of trust (Guiso, Sapienza and Zingales (2008); Gennaioli, Shleifer and Vishny (2015)); unawareness of the excess return premium (Guiso and Jappelli (2005); Grinblatt, Keloharju and Linnainmaa (2011); Cole, Paulson and Shastry (2014)); background risk (Heaton and Lucas (2000); Cocco, Gomes and Maenhout (2005)) and positive correlation of stock returns with returns of other assets in household portfolios (Benzoni, Collin-Dufresne and Goldstein (2007); Davis and Willen (2014); Bonaparte, Korniotis and Kumar (2014)); liquidity constraints, illiquid assets and consumption commitments (Grossman and Laroque (1990); Haliassos and Michaelides (2003); Chetty and Szeidl (2007)); or social interactions (Hong, Kubik and Stein (2004); Kaustia and Knüpfer (2012)).

several periods.

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 estimates 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 Nielsen Consumer Panel (NCP), the Nielsen 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 Nielsen Consumer Panel (NCP)

The [Nielsen Company Consumer Panel \(2013–2014\)](#) consists of a long-run panel of over 60,000 nationally representative households American households in 52 metropolitan areas. The goal of the NCP is to measure the detailed shopping behavior of American households while linking this data to household characteristics like household income, composition, age, and gender. 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. Nielsen 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 around \$400 of spending per month for each household.

The NCP is constructed to be a representative sample of the US population and fresh demographic information about participants is obtained each year. Nielsen maintains high quality data with regular reminders to participants that prompt them to report fully, and will remove non-compliant households from their panel. [Broda and Weinstein \(2010\)](#) provide a more detailed description of the NCP. [Einav, Leibtag and Nevo \(2010\)](#) perform a thorough analysis of the NCP, finding generally accurate coverage of household purchases though having some detectable errors in the imputed prices Nielsen uses for a subset of goods. 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. The types of goods purchased span groceries and drug products, small electronics and appliances, small home furnishings and garden equipment, kitchenware, and some soft goods. Almost all of this spending is done in-store. In our sample years, under 5% of spending in these categories is done online or via catalog purchase.

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 data. To the extent we do not observe household purchase or stockpiling of clothing, electronics, or other larger purchases, we will underestimate inventory and thus we consider our estimates a lower bound of household inventory levels.

2.2 Nielsen 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 and spans the years 2006-2014. Nielsen also provides the location of the stores at the three-digit ZIP Code level (ZIP3, e.g. 602 instead of 60208). 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 a wide range of the largest retailers in the grocery, mass merchandiser, drugstore and pharmacy, and other miscellaneous retail sectors. Within the store, the data provide a comprehensive view of products sold, with more than 2 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 mass merchandisers. In total, these data comprise over 10 billion transactions per year worth nearly \$250 billion. Here, we use data from 2013 and 2014 and merge this with the NCP to compute savings measures.

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. The SCF is a triennial survey and we use data from 2010, 2013 and 2016.

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 Nielsen 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 [National Health and Nutrition Examination Survey \(2013–2014\)](#) of the Centers for Disease Control and Prevention. 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 Nielsen.

We manually assign each of over 4,000 food items to a Nielsen product group code. The item information is very detailed, for example typically distinguishing between whether a vegetable item was canned, fresh or frozen. In some cases the NHANES code corresponds to a meal which includes multiple ingredients (for example, "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 Nielsen product module. For example, there are several codes related to fresh milk (1%, 2%, skim, whole, lactose free), but different brands or pack sizes of

products are not separated. We use data from the 2013-2014 survey.

3 Measuring Levels of Household Inventory

To compute household inventories using the NCP data we must make some assumptions. These assumptions are necessary because, although we can track the flow of purchases for different items over time in the NCP, the initial inventory and flow of consumption are not observed.

Our first assumption is that a given product’s rate of inventory depletion (rate of consumption and depreciation) is constant throughout the year, and that annual inventory depletion equals total annual spending such that inventory is stationary. The second assumption is that the initial inventory is such that inventory is never negative during each calendar year. With these two assumptions, we can compute both initial levels of household inventories as well as the inventory level at all points in time for each household.

These assumptions will not always hold at all levels of product aggregation. To be consistent with our assumption that consumption is constant throughout the year, we must aggregate individual products to a broader level. For instance, if a household switches cereal brands purchased, they are not necessarily stocking up on all brands at once, but keeping consumption of that product *group* constant. In practice, the household was consuming individual products within a category one after the other, and consumption was only continuous at a higher level of aggregation. For this reason, we combine individual products at a Nielsen “product group” level when computing inventories.⁶ We then sum over all product groups to get total household inventory.

Aggregation also makes our second assumption more conservative. 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 do not hit zero at some point during the year, our measure of inventories will be an underestimate.

Therefore, defining the product categories too narrowly can lead to inventory being overstated, but defining them too broadly can lead to inventory being understated. Using direct evidence from nutritional survey data (NHANES) on actual consumption rather than spending, Appendix A shows that consumption is indeed fairly constant when aggregating up to Nielsen product groups; [Aguiar and Hurst \(2013\)](#) provide similar evidence. Moreover, we show that even with a more conservative product aggregation choice, household inventories are still substantial and that the financial returns to household inventory management are robust within this range of inventory levels.

3.1 Computing Inventories

We next derive a formula for average inventory. The average inventory held over the period from $t = 0$ to T is $\frac{1}{T} \int_0^T I_t dt$, where I_t is the level of inventory at time t . Inventory at time t reflects the time zero level of inventory I_0 , purchases made on trips between time 0 and time t , and the rate of

⁶This still leaves us with fairly disaggregated data as Nielsen covers 118 “product groups” spanning categories such as “Crackers”, “Dough Products”, “Fresh Meat”, “Fresh Produce”, “Prepared Food Ready to Serve”, “Soft Goods”, “Automotive Products”, “Hardware and Tools”, and “Toys and Sporting Goods”. We restrict attention to products with the most common unit of measurement, “ounces.”, accounting for over half of UPCs.

inventory depletion d (i.e. consumption c and depreciation δ), which we assume to be constant:

$$I_t = I_0 + \sum_{j=1}^{n_t} S_{t_j} - d \cdot t \quad (1)$$

t_1, t_2, \dots, t_{n_t} are the dates of the n_t shopping trips occurring between time 0 and time t . S_{t_j} is the value of purchases made on the j th trip. Next, we compute the integral: $\int_0^T I_t dt = I_0 T + \sum_{j=1}^{n_T} S_{t_j} (T - t_j) - d \cdot \frac{T^2}{2}$ and we divide by T to get an expression for the average inventory:

$$\frac{1}{T} \int_0^T I_t dt = I_0 + \sum_{j=1}^{n_T} S_{t_j} \frac{(T - t_j)}{T} - d \cdot \frac{T}{2}. \quad (2)$$

When applying the formula to the data, we compute average annual inventory, so with t measured in years we have $T = 1$. Assuming annual depletion is equal to annual spending, annual average inventory of household h in product group g in calendar year y is:

$$\text{Avg. Inventory}_{h,g,y} = \text{Initial Inventory}_{h,g,y} + \sum_{j=1}^{N_{h,y}} \text{Value}_{h,g,j} \cdot \% \text{ of Year Left}_{h,j} - \frac{1}{2} \sum_{j=1}^{N_{h,y}} \text{Value}_{h,g,j}. \quad (3)$$

$N_{h,y}$ is the number of trips household h makes over calendar year y , $\text{Value}_{h,g,j}$ is the value of products purchased in product group g on shopping trip j (i.e. trip size), and “% of Year Left $_{h,j}$ ” is the share of the calendar year remaining when trip j occurs.

Our measure of inventory is 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 Section 4, we use USDA data to calibrate the product depreciation rates in our model, which allows us to explicitly take spoilage into account when computing returns.

There are various approaches we could take when computing $\text{Value}_{h,g,j}$. One approach would be to simply use household h 's total spending in product group g on trip j . However in this case fluctuations in unit prices over time may affect the inventory calculation. Instead we compute:

$$\text{Value}_{h,g,j} = P_{h,g,y} Q_{h,g,j}, \quad (4)$$

where $Q_{h,g,j}$ is the quantity purchased by household h in product group g on trip j and:

$$P_{h,g,y} = \frac{\sum_{j=1}^{N_{h,y}} \text{Spending}_{h,g,j}}{\sum_{j=1}^{N_{h,y}} Q_{h,g,j}}, \quad (5)$$

where $\text{Spending}_{h,g,j}$ is total dollar spending by household h in product group g on trip j .

We compute $\text{Initial Inventory}_{h,g,y}$ 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:

$$\text{Initial Inventory}_{h,g,y} = -P_{h,g,y} \min_j I_{h,g,t_j}. \quad (6)$$

where I_{h,g,t_j} is the inventory remaining in ounces immediately prior to trip j , assuming that $I_{h,g,0} = 0$:

$$I_{h,g,t_j} = \begin{cases} -t_j \cdot \sum_{j=1}^{N_{h,y}} Q_{h,g,j} & \text{if } j = 1 \\ I_{h,g,t_{j-1}} + Q_{h,g,j-1} - (t_j - t_{j-1}) \cdot \sum_{j=1}^{N_{h,y}} Q_{h,g,j} & \text{if } j > 1. \end{cases} \quad (7)$$

Where t_j is the time of trip j relative to the start of year. t_j is measured in years and takes values between 0 and 1. Note that $\text{Initial Inventory}_{h,g,y} \geq 0$, because $I_{h,g,t_1} \leq 0$. When computing $\text{Value}_{h,g,j}$ we restrict attention to goods measured in ounces, so $\text{Spending}_{h,g,j}$ is spending by household h on trip j on goods measured in ounces, and $Q_{h,g,j}$ is the total quantity measured in ounces. Around 50% of UPCs in the NCP are measured in ounces. The majority of the remaining products are measured in “CT” (count) – a unit which is not comparable across different UPCs and therefore inappropriate for our inventory calculation. Given this, we compute a coverage measure for each household-year:

$$\text{Coverage}_{h,y} = \frac{\text{Spending on Covered Items}_{h,y}}{\text{Total Spending}_{h,y}} \quad (8)$$

and compute the final inventory measure as:

$$\text{Inventory}_{h,y} = \frac{\sum_g \text{Avg. Inventory}_{h,g,y}}{\text{Coverage}_{h,y}} \quad (9)$$

With this approach, the average amount of household inventory in NCP goods is \$1,113.⁷ The average level of inventory varies with the level of aggregation we assume. For transparency, we compute average inventory holdings under alternative assumptions about aggregation as compared to the “Nielsen Product Group” classification that we utilize. Nielsen includes a number of levels of product classification. Aggregating to the broadest product category, the 10 “Nielsen Departments”, gives average inventory of \$716. Other possibilities include aggregating to the 1,305 “Nielsen Product Modules” or to the approximately 2 million unique product codes (UPC), i.e. not aggregating at all. This yields average inventories of \$1,379 and \$1,863, respectively. In our opinion, these latter two are likely to overstate inventories as the constant consumption assumption is probably inappropriate at such aggregation levels.

Figure I shows the distribution of our inventory measure across households. This measure of inventory naturally excludes inventory holdings in goods not covered by the NCP; most notably it excludes all large durable items like cars, furniture, most clothing and electronics. On the household

⁷We also obtain a similar value for total inventory by defining $\text{Value}_{h,g,j}$ as household h 's total spending in product group g on trip j across all units. This helps to address the concern that goods measured in ounces are not representative.

balance sheet, such 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.

Overall, inventory (even with durables excluded) is an important asset for many households. To show this, we compute household income quintiles in the SCF and use income information in the NCP to assign each Nielsen household to a quintile. We compute the average value of inventory for each household and take the median across households in each quintile during 2013 and 2014. Using data from the 2010, 2013 and 2016 SCF, we compute median financial assets within each income quintile. For each income quintile, we then compute the inventory portfolio share, $\frac{\text{Inventory}}{\text{Financial Assets} + \text{Inventory}}$. Figure II shows the inventory portfolio share by income. For households in the bottom income quintile, inventories account for around 70% of assets. As income increases, inventory holdings grow more slowly than financial assets and the inventory portfolio share declines.

In general, this section shows that inventory levels are non-trivial for many households. For a large proportion of SCF households, this liquidity need for inventory management represents a large proportion of SCF financial assets, and therefore it is plausible that this motive may explain a portion of equity market non-participation by poor households.

3.2 Validating the Household Inventory Measure

We perform two primary validation exercises for our measure of household inventory. First, we show that our measure of inventory correlates positively with measures of product life which we obtain from FSIS, a third-party source. Second, we show that for households that move, quantities purchased decrease sharply in the months leading up to the move, running down inventories. While this evidence is not conclusive, it collectively shows that our measure of inventory is associated with the product properties and household behavior we would expect.

3.2.1 Measured Inventory and Product Life

Table I shows that inventory as a share of Nielsen spending is increasing in durability. It serves as a check on the magnitudes for our calculations of inventory levels. We manually assign each Nielsen product module a usable life in months, relying primarily on the FSIS data. Product life ranges from less than a week up to five years or more. The majority of spending in the NCP is on products with a lifetime of either less than three months (27%) or more than one year (around 55%).⁸

We consider products with a lifetime of less than three months to be perishable. We define semi-perishable products as those with a lifetime of at least three months and less than one year, and products with a lifetime longer than one year are considered non-perishable. In column 1, we see that households hold about an extra 2.5 weeks of spending in semi-perishable products relative to perishables (0.048×52 weeks/year), and an extra 4.4 weeks of spending in non-perishable products relative to perishables (0.085×52 weeks/year). Columns 2 through 4 show the relationship is robust to controlling for the number of shopping trips as well as household fixed effects.

⁸In practice, product life is also household specific and depends on the way the products are prepared. For example, even perishable products such as fresh produce may be converted into a more storable form.

3.2.2 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 spending.

Figure III shows that the behavior of the subset of households who move is consistent with this story. Households cut spending well in advance of a move and by a large amount, consistent with our claim that they hold a large stockpile of inventories. Spending returns to normal immediately following the move. An alternative story 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 in-store savings decline around the move, which seems inconsistent with this story, but is consistent with households running down their stockpile. We discuss this further in Section 5.1.

Assuming that the decline in quantity purchased reflects households running down inventory, we can interpret the cumulative decline in quantity purchased 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. Similarly, if households do not plan far enough in advance, some of the stockpile might also simply be disposed of and not show up in lower purchases. Hence, the observed decline is likely much smaller than total inventory.

We find that in the sample of all households moving to a new 3-digit ZIP Code, the cumulative quantity decline is about 6.8% of annual quantity purchased.⁹ Restricting the sample to households moving more than 937 kilometers (the top quartile of move distance) the cumulative decline is about 12.3% of annual quantity purchased. This provides additional evidence, independent of our assumptions in Section 3, that households hold at least several hundred dollars of inventory.

4 A Model of Optimal Household Inventory Management

By setting aside working capital, households can reduce the average price paid for consumer products. This can act as a substitute to the channel identified by previous work that has focused on more frequent shopping trips to take advantage of lower prices (e.g. [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.

In order to understand the implications for other investment or borrowing behavior, we want to know the marginal financial return to allocating additional funds to working capital. In this section, we use the NCP to calibrate a model of optimal household inventory management. We then use the model to compute the (net) marginal returns to household working capital investment, taking into account holding costs due to depreciation and trip fixed costs, which are not directly observed in the data.

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 identify

⁹Given the limited number of households who move to a new 3-digit ZIP Code, we extend the sample period back to 2006 for this analysis.

the stockpiling channel of consumer responses to these shocks. Here, we briefly summarize the two main papers which motivate our model and identify the specific mechanisms, [Baker, Johnson and Kueng \(forthcoming\)](#) and [Hendel and Nevo \(2006a\)](#).

In [Baker et al. \(forthcoming\)](#), we use comprehensive state and local sales tax changes at a monthly frequency and show that consumers respond strongly to these changes along several margins, including stocking up on products subject to sales taxes in advance of a tax increase. This paper exploits the fact that such tax changes are always anticipated at least a couple of months in advance, giving consumers the opportunity to stock up to take advantage of the temporary low tax-inclusive prices. In a series of papers, [Hendel and Nevo \(2006a,b, 2013\)](#) identify the effect of temporary price discounts on household inventory in the context of a dynamic discrete choice model of individual product demand (e.g. laundry detergent). In addition to using various sources of identifying variation, the authors highlight the importance of household inventory to explain the large price elasticities observed in scanner data.

Our model incorporates two types of savings: buying in larger quantities (“bulk”) and buying items on sale (“deals”). This essentially drives 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 on sale for a given trip size).

We are interested in how allocating a marginal dollar to household working capital facilitates savings through each channel. The model is quite similar to [Arrow, Harris and Marschak \(1951\)](#) and the steady state version of the model in [Baker et al. \(forthcoming\)](#). The primary difference here is that households can benefit from buying in bulk and taking advantage of temporary sales, whereas the model in [Baker et al. \(forthcoming\)](#) captures intertemporal substitution behavior in response to an anticipated permanent price change induced by an anticipated consumption tax change.

Because buying large quantities reduces trip frequency and the ability to take advantage of sales, there is a trade-off between the two types of shopping policies to reduce the average unit price. In general, 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.

4.1 The Household’s Problem

The household’s problem is to minimize the cost of providing a monthly consumption flow of C , subject to an inventory constraint. For simplicity, we assume that the flow of consumption is constant both between trips and across trips.¹⁰ The cost per trip can be decomposed into two components – a fixed cost (e.g. the opportunity cost of time spent shopping) and a variable component which depends on the quantity of products purchased. The effective price per unit depends on the quantity purchased (bulk savings) and also on the household’s choice of bargain-hunting policy (deal savings).

In the model, as in the data, households consume goods with varying degrees of perishability. Allowing for varying perishability is important for matching the data and is consistent with what

¹⁰This assumption can be relaxed. For the CES case, see [Baker et al. \(forthcoming\)](#).

we observe in the NCP data. Perishable goods are important for generating a realistic trip frequency, while non-perishable goods allow us to simultaneously match the value of inventories.

The household's problem is to choose the time between trips Δ (measured in months) and bargain-hunting policies $m_l \in \mathbb{N}_0$ for each good indexed by its level of perishability, l . The policy variable m_l is the maximum number of trips in advance that a household is willing to purchase and store an item with perishability level l . All goods are purchased according to the same trip schedule determined by Δ , but households are allowed to choose a distinct value of m_l for each type of good. This is described in more detail in Section 4.5.¹¹

The household minimizes the average monthly cost of providing the exogenous consumption flow C :

$$V(\bar{I}; \theta) = \min_{\Delta, \{m_l\}} \frac{k + \sum_l P_l(\Delta, m_l) S_l(\Delta)}{\Delta} \quad (10)$$

subject to:

$$\sum_l I_l(\Delta, m_l) \leq \bar{I}. \quad (11)$$

where k is the shopping trip fixed cost, $P_l(\Delta, m_l)$ is the effective price per unit, taking into account bulk discounts, sales, and holding costs associated with setting $m_l > 0$ (i.e. stocking up in advance). $S_l(\Delta)$ is the quantity required immediately following a trip to satisfy consumption flow C until the next trip occurs.¹² Section 4.3 describes how the trip interval Δ and the trip size S_l are linked given the requirement that inventory levels neither grow without bound, nor hit zero prior to the next shopping trip. With this restriction, the household cannot choose Δ and S_l independently. Consequently, when the household chooses a trip interval, this directly implies a trip size.

The vector $\theta = (\{\delta_l, s_l\}, x, p_f, p_d, \alpha, \beta, \sigma, k, C)'$ collects the parameters of the model. δ_l are the monthly rates of depreciation for goods in each perishability group l . The effective unit price paid by the household depends on x , the probability that a particular product is on sale, as well as the full price (or "list price") p_f and the discounted price p_d . For simplicity, we assume that the sale probability is the same every trip regardless of trip length Δ . This assumption is reasonable for the typical trip intervals we observe in the Nielsen data (see the discussion below). The effective unit price paid also depends on the relationship between quantity purchased S_l and price per unit because of bulk discounts. The parameters which describe this bulk discount relationship are α , β and σ . Total consumption is $C = \sum_l C_l$, where $C_l = s_l C$ with good shares s_l such that $\sum_l s_l = 1$. The effective price function $P_l(\Delta, m_l)$ is characterized in Section 4.5.

At a given point in time, a portion of household working capital will be held as cash and the rest will be held as stored inventory goods. The inventory constraint means that the value of stored

¹¹We 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, and so households prefer to buy all goods on the same trip. For a more detailed explanation of this tradeoff, see [Bartmann and Beckmann \(1992\)](#).

¹²When calibrating the model in Section 4.8, we normalize prices such that the effective unit price equals one on average for untargeted or inattentive shopping ($m_l = 0$) and when purchasing the "standard" pack size of each product ($S_l = \hat{S}_l$ and thus $\Delta = \hat{\Delta}$), i.e. $P(\hat{\Delta}, 0) = 1$, which is therefore the price of one physical unit of S_l .

goods cannot, at any point in the shopping cycle, exceed the amount of assets set aside for managing inventory. The maximum inventory holdings occur immediately following a trip and at this point in time 100% of household working capital is held as stored inventory goods.

I_l is the level of inventory of each good with perishability l immediately following a trip (i.e. the value of inventory remaining prior to the trip, plus the value purchased $P_l S_l$) and hence $\sum_l I_l$ is equal to household working capital at this point since cash holding is zero. The level of inventory remaining immediately prior to a trip depends on x , the probability that a given product is on sale, and δ_l , its depreciation rate. This is because x and δ_l determine the household's optimal strategy for stocking up on goods when they are on sale. The more a household engages in this savings strategy, the higher the level of inventory will be when going to the store.

Ultimately, we are interested in the relationship between the dollar amount invested in household working capital and the dollar value of savings. 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 of providing the household's exogenous consumption stream so that we can define the marginal (net) return to household inventory management as $r_{\bar{I}}(\bar{I}) = V'(\bar{I})$.

While the problem does have a stochastic foundation, it is effectively deterministic. This is because we assume that, aggregating across many products, the share of products on sale each trip, x , is constant and equal for all good types l . This is fairly consistent with the NCP and NRP data, where we see regularly rotating sets of goods on sale over time, as shown in Appendix Figures B.1 and B.2.

4.2 Implications for Portfolio Choice

In the model, working capital 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. Our model focuses on the choice of shopping strategy to minimize the cost of supplying a given consumption flow. Knowing how our model fits into this higher-level problem is important for understanding the implications for stock market participation. Therefore, before solving the shopping problem we demonstrate 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. For simplicity of exposition, we assume the portfolio choice problem can be solved independently of the consumption problem.¹³ Assume 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 market portfolio. The household maximizes expected utility of end-of-

¹³The marginal return to working capital is independent of consumption assuming that both k and \hat{S}_l are scaled proportionally. This assumption is reasonable in cases where the opportunity cost of time is increasing in consumption (or "permanent income"), and where increases in consumption are reflected in purchases of higher quality products rather than purchasing larger quantities.

period wealth (or consumption) by solving the following problem:

$$\max_{\alpha_{\bar{I}}, \alpha_f, \alpha_m} EU((1 + \tilde{r}_p)w) \quad (12)$$

subject to:

$$\tilde{r}_p = \alpha_{\bar{I}}r_{\bar{I}}(\bar{I}) + \alpha_f r_f + \alpha_m \tilde{r}_m \quad (13)$$

$$\bar{I} = \alpha_{\bar{I}}w \quad (14)$$

$$1 = \alpha_{\bar{I}} + \alpha_f + \alpha_m, \quad (15)$$

where $\alpha_{\bar{I}}$ is the share of initial wealth w allocated to working capital, α_f is the share allocated to the risk-free bond with return r_f , α_m is the share allocated to the risky asset with stochastic return \tilde{r}_m , and \tilde{r}_p is the rate of return of the total portfolio.¹⁴

$r_{\bar{I}}(\bar{I})$ 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 both the amount invested ($\alpha_{\bar{I}}w$) and the level of consumption.

We treat the working capital investment as a risk-free asset. This is consistent with our assumption that after aggregating over a large number of products with independently distributed sales over a sufficiently long time period, the return is effectively deterministic.

Assuming consumers are risk averse, they choose $\alpha_{\bar{I}} = 1$ 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 $r'_{\bar{I}}(\bar{I}) > r_f$ consumers will split assets between working capital and the risky investment, as the risk-free bond is strictly dominated. As wealth becomes large, consumers will allocate all additional wealth to financial assets and the household working capital constraint (11) no longer binds. Consequently, $\alpha_{\bar{I}}$ gradually declines as wealth increases.

4.3 Quantity per Trip, $S_l(\Delta)$

Trip size $S_l(\Delta)$ is the amount of good type l which a household needs during a period of length Δ to support the constant consumption flow C_l , taking into account depreciation during that time interval $[0, \Delta]$:

$$S_l(\Delta) = \int_0^{\Delta} e^{\delta_l t} C_l dt = \frac{C_l}{\delta_l} (e^{\delta_l \Delta} - 1). \quad (16)$$

If the trip size does not satisfy this condition, either inventory will grow without bound or the household will run out of a product before the next trip.

¹⁴The average inventory portfolio share in Figure II is closely related to $\alpha_{\bar{I}}$. $\alpha_{\bar{I}}$ additionally includes a cash component, but inventory accounts for the majority of working capital in our model for realistic parameter values.

4.4 The Bargain Hunting Policy, m_l

The bargain-hunting policy m_l is the maximum number of trips in advance that a household is willing to purchase and store an item with perishability level l . Here we describe the household's strategy for taking advantage of random sales. For simplicity, we assume that the household chooses a multiple of the amount required to last over the trip interval Δ . This ensures that the household never runs out of the item. For convenience we refer to these increments as "packs".

Every trip the household must choose how much of each product to buy. Because the household faces holding costs, it does not make sense to stock up on full-priced products. However, when the household observes a product on sale, it may make sense to buy more than is required for current consumption. For example, suppose the household sees that a product is on sale, but they still have one pack left in stock (i.e. just enough inventory to provide consumption during a period of length Δ , the trip interval). This means that if they now buy additional inventory, they will need to store the product for additional time Δ before starting to consume it.

The price paid in store is p_d , but the household incurs additional holding costs which lead to the price paid being multiplied by $e^{\delta_l \Delta} \geq 1$. When the holding cost reflects product spoilage, the intuition is that the household needs to purchase $e^{\delta_l \Delta}$ units of the product now for every unit they want to have in Δ months' time. The effective price is thus $e^{\delta_l \Delta} p_d$. If the household decides not to buy the product now, the effective price is the expected price, $E[p] = x p_d + (1 - x) p_f$.¹⁵

Next, we consider what the household will do when they have two or more packs left in stock and observe that the item is on sale in store. The problem of whether to buy a j th pack n trips before running out is the same as the problem of whether to buy a $j - 1$ th pack $n + 1$ trips before running out because they both have the same effective price $e^{(n+j-1)\delta_l \Delta} p_d$.

Consequently, the household's bargain-hunting strategy can be characterized by identifying the earliest date at which they will buy a product on sale. With bargain-hunting policy m_l , the household will buy one pack m_l trips before running out, two packs $m_l - 1$ trips before, three packs $m_l - 2$ trips before, and so on. The optimal shopping strategy for deal savings is therefore completely summarized by m_l .

Note that setting $m_l > 0$ does not lead to stochastic fluctuations in trip size or effective price when aggregating across a large number of products. However, it does lead to an increase in holding costs for a given trip interval Δ , because products will be bought in advance of when they are actually required for consumption. Increases in m_l will also increase inventory I , holding Δ fixed.

4.5 The Effective Price Function, $P_l(\Delta, m_l)$

We now work out how the effective price per unit of good type l is related to the interval between household shopping trips, Δ , and the bargain-hunting policy, m_l . First, we explain how m_l affects the expected price paid in store. Intuitively, setting a high value of m_l raises the share of goods the household purchases on sale, and for large values of m_l the average price paid in store approaches the discount price p_d . We formalize this below.

¹⁵Note that some of the purchased product would also depreciate even further before it is consumed, but that additional depreciation cost applies in both cases and cancels out. Depreciation costs over the period when the pack is being consumed are captured in $S_l(\Delta)$.

4.5.1 Expected price paid in store given bargain-hunting policy m_l

Assuming that the household is fully stocked with respect to a particular product (i.e. has m_l packs currently in stock), we are interested in the probability that the next sale appears at trip $t = 0, 1, \dots, m_l$ respectively. Given that the probability of a sale is x , and sales are iid, the probability of observing t no-sale trips followed by a sale trip is $x(1-x)^t$. The probability that no sale occurs before the product runs out entirely is $(1-x)^{m_l+1}$. The probability that the item is purchased on sale, given bargain-hunting policy m_l , is therefore $\sum_{t=0}^{m_l} x(1-x)^t$. Note that this covers all possibilities since $(1-x)^{m_l+1} + \sum_{t=0}^{m_l} x(1-x)^t = 1$. The expected price paid in store given the bargain-hunting policy m_l is therefore:

$$E[p|m_l] = p_f(1-x)^{m_l+1} + p_d x \sum_{t=0}^{m_l} (1-x)^t. \quad (17)$$

The expected price with untargeted or inattentive shopping, $m_l = 0$, is $E[p|0] = xp_d + (1-x)p_f = E[p]$. As the value of m_l increases, the probability that the item is purchased on sale approaches 1 and the expected price paid approaches p_d . Hence, ignoring shopping trip fixed costs, households in the model would optimally shop continuously and buy everything on sale. In practice, prices are clearly not independently distributed when shopping occurs at a very high frequency. Upon revisiting the store an hour later, prices are likely to be unchanged. However, at shopping frequencies observed in the consumer spending data and in our calibrated model, independence is a reasonable assumption and a good approximation of the pricing pattern observed in the retail price data.

4.5.2 Adding holding costs and bulk discounts

Households incur holding costs if they stockpile items to take advantage of temporarily low prices. We model these holding costs as exponential product depreciation at rate δ_l . To properly account for these costs, we first work out how they differ across the states of the world enumerated above. Intuitively, inventories, and therefore holding costs, are lower when a sale is not observed for several trips in a row.

Goods which are purchased i trips in advance of when they are used incur additional holding costs of $e^{i\delta_l\Delta}$ relative to goods which are purchased on the trip immediately prior to consumption. When multiple packs are purchased on a given trip, each pack is stored for a different length of time before the household begins to consume it. For example, suppose the household has run out of a particular product at home and observes it on sale when they go to the store. They will buy $m_l + 1$ packs of the product (recall $m_l = 0$ corresponds to buying one pack of all products each trip). They will begin to consume one pack immediately, the second pack after time Δ , and the $m_l + 1$ st pack after time $m_l\Delta$. The total holding cost factor associated with this trip is therefore $\sum_{i=0}^{m_l} e^{i\delta_l\Delta}$.

Averaging over the $m_l + 1$ packs purchased gives $\frac{1}{m_l+1} \sum_{i=0}^{m_l} e^{i\delta_l\Delta}$ per pack. In general, if the previous sale before this trip was $t + 1$ periods ago, the holding cost factor associated with the trip is $\frac{1}{t+1} \sum_{i=0}^t e^{(m_l-i)\delta_l\Delta}$. We compute the average effective price per unit using the probabilities from

Section 4.5.1:

$$P_l(\Delta, m_l) = b(S_l) \cdot \left[p_f(1-x)^{m_l+1} + p_d x \sum_{t=0}^{m_l} (1-x)^t \frac{1}{t+1} \sum_{i=0}^t e^{(m_l-i)\delta_l \Delta} \right]. \quad (18)$$

For simplicity, we assume that the bulk discount function $b(S_l)$ is applied directly to the trip size $S_l(\Delta)$. Households can increase S_l , and take advantage of bulk discounts, by shopping less frequently (see Section 4.3). Bulk discounts therefore tend to raise the trip interval Δ . We specify the bulk price discount function b to match bulk discounts observed in NRP data using the following functional form, implying that unit prices decline as the quantity purchased per trip S_l increases:

$$b(S_l) = \alpha + \beta e^{-\sigma \frac{S_l}{\hat{S}_l}}. \quad (19)$$

\hat{S}_l is the trip size associated with purchasing standard packs of each item in the NRP and we will calibrate (α, β, σ) such that $b(\hat{S}_l) = 1$.

The 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 price in the model to 1 when purchasing the standard pack size ($S_l = \hat{S}_l$) and in the absence of targeted deal shopping ($m_l = 0$). This means that α is interpreted as one minus the maximum % savings which can be obtained from buying in bulk.

To set \hat{S}_l , we solve the model without bulk discounts (i.e. with $b(S_l) = 1$) and compute the optimal trip interval $\hat{\Delta}$. We then set $\hat{S}_l = S(\hat{\Delta})$. Because we calibrate the price distribution so that $E[p] = p_f x + p_d(1-x) = 1$, the expected price per unit of S_l in the model is normalized to 1 for households purchasing the standard pack size ($S_l = \hat{S}_l$) and using an untargeted shopping strategy ($m_l = 0$), i.e. $P_l(\hat{\Delta}, 0) = 1$. The calibration is described in detail in Section 4.8.

4.6 The Household Working Capital Constraint

We work out how much inventory is left over at trip time in order to test whether the household working capital constraint is satisfied, $\sum_l I_l(\Delta, m_l) \leq \bar{I}$. We value inventory at its total effective price. Only goods for which there was a sale in the previous m_l trips are still in stock immediately prior to a trip. The amount left in stock depends on how long ago the most recent sale was. If the most recent sale occurred on the previous trip, there will still be m_l packs left in stock. The total effective price per pack is $p_d \frac{\sum_{i=1}^{m_l} e^{i\delta_l \Delta}}{m_l}$. In general, if the last sale occurred $t+1$ trips ago, the value per pack of inventory in stock prior to the current trip is $p_d \frac{\sum_{i=1}^{m_l-t} e^{(i+t)\delta_l \Delta}}{m_l-t}$.

The share of goods for which the most recent sale occurred $t+1$ trips ago is $x(1-x)^t$, i.e. the probability of a sale event followed by t non-sale events. Immediately following each trip, the value of inventory of good type l is therefore:

$$I(\Delta, m_l) = P_l(\Delta, m_l) S_l(\Delta) + \mathbb{1}_{\{m_l > 0\}} S_l(\Delta) p_d \sum_{t=0}^{m_l-1} x(1-x)^t \sum_{i=1}^{m_l-t} e^{(i+t)\delta_l \Delta}. \quad (20)$$

That is, the expenditure on the trip, $P_l S_l$, plus the value of inventory accumulated on previous trips to be consumed after the current trip. If $m_l = 0$, inventory is just the current trip value. In this

case, inventory hits zero immediately prior to the next trip, and average inventory is $P_l S_l / 2$. The probability mass $\sum_{t=0}^{m_l-1} x(1-x)^t$ is equal to $1 - (1-x)^{m_l}$, where $(1-x)^{m_l}$ is the share of goods for which no inventory remains at trip time.

4.7 Solution Method

We start by defining a grid over trip intervals Δ and bargain-hunting strategies $\{m_l\}$. We then search over all combinations for which the household working capital constraint is satisfied and find the combination that minimizes the cost function.

1. Define a grid over trip intervals Δ and bargain-hunting strategies $\{m_l\}$, where $m_l \in \mathbb{N}_0$.
2. For each possible combination of $(\Delta, \{m_l\})$, compute $S_l(\Delta)$, $I_l(\Delta, m_l)$, and the value of the cost function, $\frac{k + \sum_l P_l(\Delta, m_l) S_l(\Delta)}{\Delta}$.
3. Find the values of $(\Delta^*, \{m_l^*\})$ that minimize the cost function subject to the household working capital constraint.

4.8 Calibration

We calibrate the model by choosing θ to match a number of data moments summarized in Table II. We set the fixed cost per shopping trip to $k = \$4.85$ as in Baker et al. (forthcoming). As discussed there, this is consistent with a hourly reservation wage of \$7.12 to \$10.69.

To calibrate p_f and p_d , we estimate the average price drop associated with a discount event in the NCP.¹⁶ Given an estimated log price difference of 0.293, we set $\frac{p_f}{p_d} = e^{0.293} = 1.34$. We pin down p_f and p_d by normalizing the expected price achieved using an untargeted shopping strategy without bulk discount, $E[p] = xp_d + (1-x)p_f = 1$.

To calibrate the bulk discount function $b(S_l)$, we first choose the value \hat{S}_l which corresponds to the standard trip size for which there is no bulk price discount, i.e. $b(\hat{S}_l) = 1$. We set $\hat{S}_l = S_l(\hat{\Delta}) = \frac{C_l}{\delta_l}(e^{\delta_l \hat{\Delta}} - 1)$, where $\hat{\Delta}$ is the optimal trip interval in the model without bulk discounts.¹⁷ Because we also set $E[p] = 1$, this means that the expected effective price per unit of S_l in the model is normalized to 1 for households with the standard trip size using an untargeted shopping strategy, i.e. $P(\hat{\Delta}, 0) = E[p] = 1$.

We calibrate the parameters of the function $b(S_l)$ by estimating the following relationship with weighted least squares:

$$\text{Price}_{p,q} = a_0 + a_1 e^{-\sigma \text{Units}_{p,q}}, \quad (21)$$

where $\text{Price}_{p,q}$ is the standardized unit price in pack size quintile q for product p and $\text{Units}_{p,q}$ is the standardized number of units of product p in a pack in the q th quintile.

¹⁶We define a discount indicator $D_{u,t}$ which is equal to 1 when product u in month t is purchased either with a coupon, or at a price which is lower than than the annual UPC-ZIP3 average price. We estimate the average log discount by running the regression $\log \text{Price}_{u,t} = \beta_1 D_{u,t} + \beta_2 \text{Bulk}_{u,t} + \text{Month FE}_t + \text{Product}_{u,t} \times \text{Household FE}_i + \epsilon_{u,t}$, where u indexes transactions at the UPC level.

¹⁷Our assumption is that the standard trip size corresponds to the quantity households in the model would want to buy if bulk discounts did not exist.

Next we describe how we construct $\text{Price}_{p,q}$ and $\text{Units}_{p,q}$. First, we prepare the Nielsen data by creating a new product ID. Because each pack size has a unique UPC, we need to create a broader product definition to examine the relationship between unit price and pack size holding the product fixed. Ideally we want to group otherwise identical products which are available in different sized packages. Our approach is to group products based on product module, brand, and common consumer name.¹⁸

We then compute pack size quintiles for each product that exists in multiple sizes based on total ounces in the pack. To calibrate relationship (19) we want to express both prices and pack size relative to a standard pack size for that product. To do this, we compute the average number of units in the second quintile of pack sizes for each product, as well as the expenditure weighted annual average price per unit. We then divide by these second quintile averages. That is, we assume the second quintile of pack size corresponds to the "standard" trip size in the model. That is, households with the standard trip size purchase pack sizes in the second quintile for all products. This is motivated by the observation that the first quintile appears to contain travel size packs which are substantially more expensive on a per unit basis, yielding implausibly large bulk discounts.

In the model, we assume that households purchase the same multiple of the standard pack size across all the products with the same perishability. When households in the model take advantage of bulk discounts by increasing trip size, this necessarily coincides with reduced trip frequency (holding consumption fixed).

In the data, a large proportion of spending is accounted for by products which have only limited bulk savings potential (for example, it may not be possible to purchase a pack size more than 1.5 times the standard pack size). This means potential bulk savings would be overstated if we were to estimate the relationship between pack size and price without further adjustments, as the relationship at higher pack sizes would be based only on products for which extreme pack sizes are available. For the purpose of calibrating the model, we are interested in the bulk savings the household would achieve if they increased the pack size uniformly across products consumed.

We therefore aggregate the data to pack size groups \times product, and then make sure the dataset is balanced (i.e. every product has a non-missing price for each pack size group). For products where large pack sizes are not available, we fill in the unit price for larger-than-feasible pack size groups with the unit price associated with the largest available pack size. We compute total expenditure for each product and use this to weight our regressions.

We then estimate equation (21) for different values of σ and choose σ to maximize the within- R^2 . This yields $\hat{\sigma} = 1.75$ and the estimates of a_0 and a_1 from the same specification are $\hat{a}_0 = 0.82$ and $\hat{a}_1 = 1.18$ respectively. For 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 calibrate α and β using $\alpha = \frac{\hat{a}_0}{\hat{a}_0 + \hat{a}_1 e^{-\hat{\sigma}}}$ and $\beta = \frac{\hat{a}_1}{\hat{a}_0 + \hat{a}_1 e^{-\hat{\sigma}}}$. Figure IV compares the bulk discount function we use in the model with the corresponding relationship in the data.

¹⁸We manually inspect and drop combinations where this approach is problematic because the group is likely to contain products which are not identical. For example, we drop store brands because this group contains a large number of products that are likely to be different from each other. We also drop video products and nail polish – these modules contain a large number of products that are not easily substitutable because they are typically different colors or different films.

To calibrate the probability of a sale, x , we use the deal flag from the NCP, which is equal to one for purchases where a coupon was used, or where the household considered the item to be on sale. The calibrated value of x is 24%. This means the discounted price is observed about every 4th trip. We estimate the following relationship using the NCP:

$$\text{Deal \& Coupon Share}_{h,y} = \alpha + \beta_1 \text{Inventory Ratio}_{h,y} + \beta_2 \text{Trips}_{h,y} + \epsilon_{h,y}. \quad (22)$$

Deal & Coupon Share $_{h,y}$ is the share of items purchased by household h in calendar year y where either a coupon was used or the household perceived the item to be on sale. Inventory Ratio $_{h,y}$ is household h 's average inventory in year y (computed in Section 3) divided by household h 's total Nielsen spending in year y . We then set x equal to the predicted deal and coupon share for a household with an inventory ratio equal to $0.5/\text{Trips}_{h,y}$, where Trips $_{h,y}$ is the number of trips made by household h in year y . This is the inventory ratio a household would have in the model if they did not engage in strategic deal shopping.

We calibrate perishable depreciation δ_1 and its share s_1 by combining USDA information on product storage from FSIS with information on households' purchases from the NCP. We define perishable goods as products which have expiration dates of less than one month. Using the FSIS data, we manually assign a time to expiry for each Nielsen product module.

We assume exponential product depreciation and calibrate δ_l such that the average expiry time, $\frac{1}{\delta_l}$, matches the expiry time given by FSIS for the relevant set of products. The average time to expiry for perishable goods is 0.347 months, giving $\delta_1 = 2.88$. Given an average expenditure share of non-perishable goods in the NCP of 76.9%, we set $s_1 = 0.769$.¹⁹

Finally, we calibrate C and δ_0 jointly to match the average monthly spending and inventory in the NCP. The average monthly spending on goods covered by Nielsen is \$392, the average level of household inventory is \$1,113. Because it is not in general possible to match both quantities exactly, we use values of the parameters that minimize the sum of squared percentage deviations between the model and the data. The corresponding model values are monthly spending of \$400 and average inventory of \$1,092. As discussed above, the quantity units in the model are normalized so that the price per unit in the absence of any deals or bulk discounts is \$1. The calibrated value of C is 492 units.

5 Financial Net Returns to Household Inventory Investment

Solving the optimization problem (10) yields the average monthly cost $V(\bar{I})$ of supplying consumption flow C . 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 not just the price paid in store but also shopping trip fixed costs and depreciation costs. In practice, the cost function is not smooth because m_l is discrete. Consequently, the marginal return $\frac{V(\bar{I}_0) - V(\bar{I}_1)}{\bar{I}_1 - \bar{I}_0}$ may be zero when $\bar{I}_1 - \bar{I}_0$ is small, but substantial when the increment is increased.

¹⁹One way to think about this assumption is that the expiry time for a product l has a Poisson distribution with parameter δ_l . While this means there is some chance that items will last a lot longer than the average expiry time, it does still capture the high cost of storing perishable items and generates a trip frequency consistent with the data.

It therefore makes sense to consider a somewhat larger increment. In the tables below we use an increment of 5% of annual spending. We multiply by 12 to convert monthly to annual returns.²⁰

Table III shows how increasing the maximum household working capital \bar{I} affects inventory and cash holdings, as well as the different sources of savings households are able to achieve. When the amount of funds allocated to household working capital is low, the household is restricted in its ability to take advantage of deals. This is because stockpiling products well in advance of when they are needed (i.e. a large m_l) is working capital intensive. Low levels of household working capital investment therefore constrain households to choose a low value for m_l (conditional on the fixed trip cost being non-trivial).

As the household working capital investment is increased, households choose progressively higher values of m_0^* .²¹ Under a deal-focused strategy households tend to shop more frequently, which reduces their trip size and tends to weigh on their bulk savings. With a deal-focused shopping strategy households make small and frequent trips on which they buy only a subset of their consumption bundle. They purchase only goods that are on sale, or goods for which inventory has been run down to zero.

At the same time, an increase in working capital also allows households to spend more per trip, increasing the trip interval and reducing trip fixed costs all else constant. This force works in the opposite direction, pushing the trip interval up and raising bulk savings. Given the parameter values we use here, increasing the trip interval does not exert much downward pressure on m_0^* and deal savings. This is because the non-perishable good has very low holding costs, and the perishable good is not worth stockpiling at parameters we consider. Consequently, we see the optimal trip interval increase slightly with working capital.

In general, model outcomes such as the marginal return, minimum inventory, trip interval and percentage savings need not be monotonic in \bar{I} . As discussed above, the cost function is non-smooth because m_l is discrete. Furthermore, relaxing the constraint may have positive or negative effects on inventory and savings of each type. The household may use additional working capital to increase bulk savings and reduce fixed costs, or it may use it to increase deal savings. 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 minimum inventory. Alternatively, if the household uses the additional funds to buy items on sale well in advance, this tends to put downward pressure on trip size due to depreciation costs and reduces bulk savings.

At low levels of household working capital investment, the marginal return to additional investment is very high. We compute marginal returns using a working capital increment of 5% of annual spending. When household working capital is equal to 5% of annual spending, the marginal return is 82%. The marginal return gradually diminishes and reaches zero when household working capital is around three times monthly spending.

²⁰We assume that the working capital investment remains fixed at \bar{I} throughout the year. Given that the marginal return is diminishing in \bar{I} , it is not appropriate to assume the proceeds can be reinvested at the same rate of return. To the extent that monthly returns are invested elsewhere and earn a positive return, the annual return will be larger than what we assume.

²¹Because of the high depreciation costs associated with stockpiling perishable items, m_1^* is always equal to zero for the parameter values we consider here.

We obtain qualitatively similar results when using a more conservative value of inventory to calibrate the model. Table B.1 shows the results using an average inventory level of \$716. As discussed in Section 3.1, this is the average value of inventories we obtain when aggregating to Nielsen departments rather than product groups when computing inventories. Table B.2 shows the results using inventory of \$1,379 obtained when aggregating to Nielsen product modules.

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 (optimal inventory management) and leaves out other motives such as precautionary liquidity or speculative motives to take advantage to temporary investment opportunities. Furthermore, our model applies to other goods not covered by the Nielsen data which also require additional cash holdings.

In Table IV we compute returns for households who have a very large trip fixed cost of \$100. Such high trip costs could be relevant for households with very high wages and opportunity costs or could be relevant in periods where going to the store may incur non-financial costs such as risk of disease transmission. High trip fixed costs increase the returns to working capital. With a longer interval between trips, households need a substantial amount of working capital to cover the high in-store cost associated with large trips, as well as stock up on non-perishable goods in response to deals. Table IV shows that at low levels of working capital households devote their resources to covering the cost of large trips, and forgo deal savings. At higher levels of working capital, households can afford to both maintain a large trip size and take advantage of deal savings. The fact that the perishable good share in our model is fixed at normal levels restricts the extent to which households can increase the trip interval. In practice, allowing for substitution away from these perishable products when trip fixed costs rise would allow for even larger reductions in trip frequency as the fixed cost increases (and larger returns to working capital).

5.1 Testable Implications for In-Store Savings

In the model, households obtain high returns to working capital by exploiting temporary sales and bulk purchases. The model generates a number of predictions for the relationships between working capital \bar{I} , the trip fixed cost k , and in-store savings which we can test using the NCP. This helps support our claim that financial returns to working capital are an important determinant of households' high inventory holdings in practice.

When testing model predictions using the NCP, there are two modifications we need to make given data limitations. Firstly, we use our inventory measure from Section 3 in the place of working capital, because we do not observe cash directly in the NCP.²² Secondly, we use the average number of days between a household's trips as a proxy for their trip fixed cost.

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 savings and the trip interval. In the second exercise, we vary working capital \bar{I} and plot the relationship between savings and average

²²In the model, changes in household working capital are closely related to changes in average inventory. Conditional on holding a modest amount of household working capital, the model predicts that cash holdings display little relationship to the total amount allocated to household working capital. Instead, the additional household working capital is reflected in higher inventory holdings.

inventory generated by this variation. To compare these relationships with the data, we first need to construct measures of deal savings and bulk savings using the NCP.

5.1.1 Measuring Base Prices for Goods

We compute a base price in the data which corresponds to the expected price in the model when $m = 0$ and $b(S) = 1$. That is, the average price paid for a product if the household engages in “untargeted shopping” (or “inattentive shopping”) in their area and buys the “standard” pack size. This ensures that our data definitions of savings match the model definitions as closely as possible. For this we require an alternative data source, as the NCP only provides us with the price the household actually paid for the item, not the prices which were available to them. To compute the base price we therefore use the Nielsen Retailer Panel (NRP), which provides weekly UPC price data at the store level for all products. For product p sold in 3-digit Zip Code z in calendar year y , the base price is computed using the following formula:

$$\text{AvgSecondQuintilePrice}_{z,p,y} = \frac{1}{|U_p|} \sum_{u \in U_p} \text{AvgPrice}_{z,u,y}, \quad (23)$$

where U_p is the set of UPCs associated with the second pack-size quintile of product p , and product p corresponds to the set of UPCs with the same product module, brand and common consumer name, as described in Section 4.8. $\text{AvgPrice}_{z,u,y}$ is the average price at which UPC u was sold in year y in 3-digit Zip Code z :

$$\text{AvgPrice}_{z,u,y} = \frac{\sum_{w \in W_y} \sum_{s \in S_{z,w}} P_{u,w,s} 1[Q_{u,w,s} > 0]}{\sum_{w \in W_y} \sum_{s \in S_{z,w}} 1[Q_{u,w,s} > 0]}, \quad (24)$$

where W_y is the set of weeks in calendar year y , $S_{z,w}$ is the set of stores in the NRP located in 3-digit Zip Code z in week w , $Q_{u,w,s}$ is the quantity of UPC u sold by store s in week w , and $P_{u,w,s}$ is the average per unit price at which UPC u is sold in week w by store s . $P_{u,w,s}$ is not observed for UPC-week-store combinations with zero sales ($Q_{u,w,s} = 0$).

5.1.2 Deal Savings

We compute annual deal savings using the following formula:

$$\% \text{Deal savings}_{h,y} = \frac{\sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t} - \text{Spending}_{h,y}}{\sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}}, \quad (25)$$

where $T_{h,y}$ is the set of trips taken by household h in calendar year y . $\text{Spending}_{h,y}$ is the amount that household h spent in year y . This reflects any savings from coupons.²³ $\sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}$ is the amount that household h would have spent on trip t if they paid the average UPC-ZIP3 price over that year.

That is, deal savings represents the difference between the average unit price of a UPC and the price the household actually pays: the *additional* savings that resulting from strategic shopping be-

²³Using the variable names from the NCP documentation, our measure of spending is defined as “total_price_paid” less “coupon_value”, where “total_price_paid” is the total price paid before coupon discounts, and “coupon_value” is the value of coupon discounts.

havior relative to random shopping (either over time or across stores in the same 3-digit Zip Code). Note that because different pack sizes of the same product have different UPCs, this measure does not incorporate bulk savings.

Figure Va plots the model relationship between deal savings and the trip interval generated by varying k . Deal savings are lower for households with high fixed costs. Figure Vb shows that deal savings are also lower for households with longer trip intervals in the NCP and the estimated coefficient is broadly similar. This is consistent with Aguiar and Hurst (2007)'s observation that households in the NCP who shop more frequently obtain higher savings. In our model, shopping more frequently allows households to set a higher value of m , holding working capital fixed, and therefore obtain higher savings because prices are observed more frequently.

Next, we vary working capital \bar{I} and plot the relationship between the inventory ratio and savings. Because the trip interval is a key driver of variation in the inventory ratio in the data, we control for the trip interval in this analysis. This is important, as the model implies, firstly, a positive relationship between the trip cost k and inventory, and secondly, a negative relationship between trip cost k and deal savings, all else constant. Consequently, variation in trip costs generates a negative relationship between inventory and deal savings. To measure the effect of other sources of variation in inventory, such as working capital, it is therefore essential to control for trip costs.

Figure Vc shows that increasing working capital allows households in the model to obtain more savings at a given trip frequency. This is because additional working capital allows households to set a higher value of m , holding trip frequency fixed, which increases their deal savings. Figure Vd shows that we observe a similar relationship between inventory and deal savings in the NCP. Working capital 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. high k). This is because 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 only a small level of working capital because they shop very frequently.

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, product preferences, or in the way grocery items are processed and stored. Variation in holding costs generates a similar relationship between inventory and deal savings as variation in working capital in the model. Figure Vd is therefore also consistent with variation in inventory in the data reflecting variation in holding costs.

As another test of our model's predictions, we study the savings obtained by households around the time they move. Because it is costly to move a stockpile of household goods, we expect that households will reduce m in advance of a move. That is, close to a move households may be hesitant about stocking up in response to temporary sales. Informally, we expect this will lead to a reduction in percentage savings obtained in the weeks prior to a move. 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. Figure VII shows that deal savings drop in the months prior to the move and recover gradually in the months following the move.

5.1.3 Bulk Savings

We compute bulk savings in the data using the following formula:

$$\% \text{Bulk savings}_{h,y} = \frac{\sum_t \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t} - \sum_t \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}}{\sum_{t \in T_{h,y}} \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t}}, \quad (26)$$

where index t runs through the set $T_{h,y}$ of trips taken by household h in calendar year y , and $\text{AvgSecondQuintilePrice}_{z,p,y}$ and $\text{AvgPrice}_{z,u,y}$ are defined in Section 5.1.1, $Q_{p,h,t}$ is the quantity of product p purchased by household h on trip t and $Q_{u,h,t}$ is the quantity of UPC u purchased by household h on trip t . That is, the measure compares the average price of the UPC actually purchased with the average price of a second quintile pack size of the same product.

Figure VIa plots the model relationship between trip interval and bulk savings. In the model, households with high fixed costs, and therefore less frequent trips, obtain more bulk savings by assumption. This is because households with less frequent trips also buy more each trip, holding consumption fixed. We see a similar relationship, though much more muted, in the data (Figure VIb).

The flat relationship between the inventory ratio and bulk savings in the model also follows from our assumption that bulk savings are obtained by choosing a larger steady state trip size (Figure VIc). Consequently, any 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 described in Section 5.1.2. Figure VI d shows the relationship in the data is also basically flat, supporting our modelling assumptions. That is, in both the data and the model, stockpiling reflects households taking advantage of temporary deals on standard pack sizes, rather than buying in bulk. One possible reason for this is that, for some products, depreciation is much higher after the pack is opened. Therefore, for a household who wishes to continuously consume a particular product, depreciation costs will be lower if the household stockpiles many small packs of the item, rather than a small number of large packs.

5.1.4 Total Savings

We compute total percentage savings as the sum of dollars of deal and bulk savings divided by total spending,

$$\begin{aligned} \% \text{Total savings}_{h,y} &= \frac{\sum_{t \in T_{h,y}} \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t} - \sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}}{\text{Spending}_{h,y}} \\ &+ \frac{\sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t} - \text{Spending}_{h,y}}{\text{Spending}_{h,y}} \end{aligned}$$

which can be simplified to:

$$\% \text{Total savings}_{h,y} = \frac{\sum_{t \in T_{h,y}} \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t} - \text{Spending}_{h,y}}{\text{Spending}_{h,y}}. \quad (27)$$

Previously, we computed the *net* marginal return to working capital using the calibrated model.

An alternative measure of returns which, in principle, can be computed using the data alone is the “gross” return, which ignores fixed costs and holding costs. This measure of returns reflects in-store savings only. The gross return associated with increasing working capital from \bar{I}_0 to \bar{I}_1 is therefore:

$$\text{Gross return} = \frac{\text{Dollar in-store savings}(\bar{I}_1) - \text{Dollar in-store savings}(\bar{I}_0)}{\bar{I}_1 - \bar{I}_0}. \quad (28)$$

Unfortunately, we do not directly observe working capital in the data and we also do not have an exogenous source of variation. However, we can construct the theoretical relationship using the model and compare this with the data.

We vary working capital \bar{I} in the model and compute the level of both in-store savings and inventory. We then divide both variables by total spending and plot the relationship. We also plot the relationship in the NCP between total savings computed using equation (27) and the ratio of inventory to total spending.

Figure VIIIa shows that varying working capital generates a positive relationship between savings and inventory in the model, holding the trip interval fixed. Figure VIIIb shows that there is also an increasing relationship in the NCP.²⁴ While we control for the number of trips (as a measure of fixed costs), the NCP results are only suggestive because many other factors could be driving the inventory variation.

Earlier in the paper we analyzed the behavior of households who move. As moving a stockpile of household goods a long distance is costly, we expect that the net return to working capital (including moving costs) declines for these households as the move date approaches. Consequently, this should lead to a reduction in the level of working capital \bar{I} optimally chosen by these households. This generates variation in \bar{I} we can in principle use to measure the relationship between \bar{I} and in-store savings in the data.

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.

Our findings are highly relevant for understanding the ability of households to support consumption smoothing after shocks to income and spending. We demonstrate that households earn high returns from inventory management through several channels at low levels of inventory, but these 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 it can even dominate some forms of borrowing costs such as credit card interest rates and fees.

²⁴Appendix Figure B.3 shows that similar results obtain when using different levels of product aggregation to compute average household inventories rather than product groups.

References

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino**, "Perception of House Price Risk and Homeownership," *NBER Working Paper No. 25090*, 2020.
- Aguiar, Mark and Erik Hurst**, "Life-Cycle Prices and Production," *American Economic Review*, 2007, 97 (5), 1533–1559.
- and —, "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 2013, 121 (3), 437–492.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg**, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, 15 (3), 47–68.
- Arrow, Kenneth J., Theodore Harris, and Jacob Marschak**, "Optimal Inventory Policy," *Econometrica*, 1951, pp. 250–272.
- Baker, Scott R., Stephanie Johnson, and Lorenz Kueng**, "Shopping for Lower Sales Tax Rates," *American Economic Journal: Macroeconomics*, forthcoming.
- Barberis, Nicholas, Ming Huang, and Richard H. Thaler**, "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," *American Economic Review*, 2006, 96 (4), 1069–1090.
- Bartmann, Dieter and Martin J. Beckmann**, *Inventory Control: Models and Methods*, Springer-Verlag, 1992.
- Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein**, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, 2007, 62 (5), 2123–2167.
- Bertaut, Carol C. and Martha Starr**, "Household Portfolios in the United States," *FEDS Working Paper No. 2000-26*, 2000.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, "Behavioral Household Finance," in Douglas B. Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics: Applications and Foundations*, Vol. 1, Elsevier, 2018, pp. 177–276.
- Black, Sandra E., Paul J. Devereux, Petter Lundborg, and Kaveh Majlesi**, "Learning to Take Risks? The Effect of Education on Risk-Taking in Financial Markets," *Review of Finance*, 2018, 22 (3), 951–975.
- Boizot, Christine, Jean-Marc Robin, and Michael Visser**, "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, April 2001, 111 (470), 391–419.
- Bonaparte, Yosef, George M Korniotis, and Alok Kumar**, "Income Hedging and Portfolio Decisions," *Journal of Financial Economics*, 2014, 113 (2), 300–324.
- Broda, Christian and David E. Weinstein**, "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 2010, 100 (3), 691–723.
- Chetty, Raj and Adam Szeidl**, "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, 2007, 122 (2), 831–877.


- Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi**, "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 2003, 93 (1), 15–37.
- Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout**, "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 2005, 18 (2), 491–533.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong**, "The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review*, 2015, 105 (3), 993–1029.
- Cole, Shawn, Anna Paulson, and Gauri Kartini Shastry**, "Smart Money? The Effect of Education on Financial Outcomes," *Review of Financial Studies*, 2014, 27 (7), 2022–2051.
- Davis, Steven J. and Paul. Willen**, "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice," *Quarterly Journal of Finance*, 2014, 3, 1–53.
- Einav, Liran, Ephraim Leibtag, and Aviv Nevo**, "Recording Discrepancies in Nielsen Homescan Data: Are they Present and Do They Matter?," *QME*, 2010, 8 (2), 207–239.
- Epstein, Larry G. and Martin Schneider**, "Ambiguity and Asset Markets," *Annual Review of Financial Economics*, 2010, 2 (1), 315–346.
- Fisman, Raymond and Inessa Love**, "Trade Credit, Financial Intermediary Development, and Industry Growth," *Journal of Finance*, 2003, 58 (1), 353–374.
- Food Safety and Inspection Service FoodKeeper Data**, United States Department of Agriculture 2020. Available at <https://catalog.data.gov/dataset/fsis-foodkeeper-data> (accessed July 7, 2020).
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny**, "Money Doctors," *Journal of Finance*, 2015, 70 (1), 91–114.
- Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo**, "Consumer Shopping Behavior: How Much do Consumers Save?," *Journal of Economic Perspectives*, 2009, 23 (2), 99–120.
- Grinblatt, Mark, Matti Keloharju, and Juhani Linnainmaa**, "IQ and Stock Market Participation," *Journal of Finance*, 2011, 66 (6), 2121–2164.
- Grossman, Sanford and Guy Laroque**, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 1990, 58 (1), 25–51.
- Guiso, Luigi and Paolo Sodini**, "Household Finance: An Emerging Field," in George M. Constantinides, Milton Harris, and Rene M. Stulz, eds., *Handbook of the Economics of Finance*, Vol. 2, Elsevier, 2013, pp. 1397–1532.
- **and Tullio Jappelli**, "Awareness and Stock Market Participation," *Review of Finance*, 2005, 9 (4), 537–567.
- **, Paola Sapienza, and Luigi Zingales**, "Trusting the Stock Market," *Journal of Finance*, 2008, 63 (6), 2557–2600.
- Haliassos, Michael and Alexander Michaelides**, "Portfolio Choice and Liquidity Constraints," *International Economic Review*, 2003, 44 (1), 143–177.
- **and Carol C. Bertaut**, "Why do so Few Hold Stocks?," *Economic Journal*, 1995, 105 (432), 1110–1129.
- Heaton, John and Deborah Lucas**, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 2000, 55 (3), 1163–1198.

- Hendel, Igal and Aviv Nevo**, “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 2006, 74 (6), 1637–1673.
- **and** —, “Sales and Consumer Inventory,” *RAND Journal of Economics*, 2006, 37 (3), 543–561.
- **and** —, “Intertemporal Price Discrimination in Storable Goods Markets,” *American Economic Review*, 2013, 103 (7), 2722–2751.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein**, “Social Interaction and Stock-Market Participation,” *Journal of Finance*, 2004, 59 (1), 137–163.
- Hurd, Michael, Maarten Van Rooij, and Joachim Winter**, “Stock Market Expectations of Dutch Households,” *Journal of Applied Econometrics*, 2011, 26 (3), 416–436.
- Jappelli, Tullio and Luigi Pistaferri**, “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2010, 2.
- , **Luigi Guiso, and Michael Haliassos**, *Household Portfolios*, MIT Press, 2002.
- Kaplan, Greg and Guido Menzio**, “Shopping Externalities and Self-Fulfilling Unemployment Fluctuations,” *Journal of Political Economy*, 2016, 124 (3), 771–825.
- **and Sam Schulhofer-Wohl**, “Inflation at the Household Level,” *Journal of Monetary Economics*, 2017, 91 (C), 19–38.
- Kaustia, Markku and Samuli Knüpfer**, “Peer Performance and Stock Market Entry,” *Journal of Financial Economics*, 2012, 104 (2), 321–338.
- Kézdi, Gábor and Robert J. Willis**, “Stock Market Expectations and Portfolio Choice of American Households,” *University of Michigan Working Paper*, 2009.
- Kueng, Lorenz**, “Excess Sensitivity of High-Income Consumers,” *Quarterly Journal of Economics*, 2018, 133 (4), 1693–1751.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman**, “A Debt Puzzle,” in Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press, 2003, chapter 11, pp. 228–266.
- Luttmer, Erzo G.J.**, “What Level of Fixed Costs Can Reconcile Consumption and Stock Returns?,” *Journal of Political Economy*, 1999, 107 (5), 969–997.
- Malmendier, Ulrike and Stefan Nagel**, “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” *Quarterly Journal of Economics*, 2011, 126 (1), 373–416.
- Mankiw, N. Gregory and Stephen P. Zeldes**, “The Consumption of Stockholders and Non-Stockholders,” *Journal of Financial Economics*, 1991, 29 (1), 97–112.
- National Health and Nutrition Examination Survey**, Centers for Disease Control and Prevention 2013–2014. Available at <https://www.cdc.gov/nchs/nhanes/index.htm> (accessed September 17, 2020).
- Nevo, Aviv and Arlene Wong**, “The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession,” *International Economic Review*, 2019, 60 (1), 25–51.
- Nielsen Company Consumer Panel**, Kilts Center for Marketing at the University of Chicago Booth School of Business 2013–2014. Available at <https://www.chicagobooth.edu/research/kilts/datasets/nielsen> (accessed December 2, 2015).

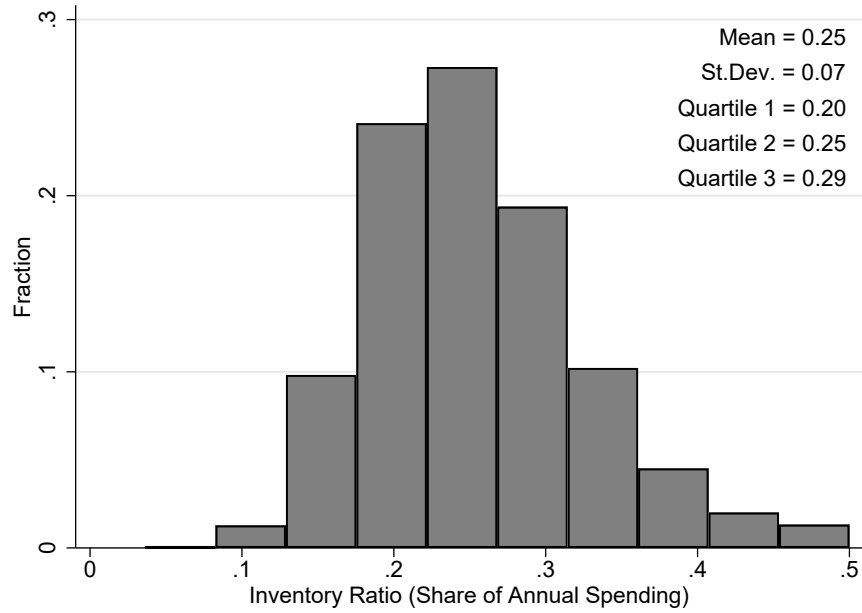
- Nielsen Company Retail MSR Scanner Data**, Kilts Center for Marketing at the University of Chicago Booth School of Business 2013–2014. Available at <https://www.chicagobooth.edu/research/kilts/datasets/nielsen> (accessed December 2, 2015).
- Petersen, Mitchell A. and Raghuram G. Rajan**, “Trade Credit: Theories and Evidence,” *Review of Financial Studies*, 1997, 10 (3), 661–691.
- Rooij, Maarten Van, Annamaria Lusardi, and Rob Alessie**, “Financial Literacy and Stock Market Participation,” *Journal of Financial Economics*, 2011, 101 (2), 449–472.
- Samphantharak, Krislert and Robert M. Townsend**, *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting*, Cambridge University Press, 2010.
- Stroebel, Johannes and Joseph Vavra**, “House Prices, Local Demand, and Retail Prices,” *Journal of Political Economy*, 2019, 127 (3), 1391–1436.
- Survey of Consumer Finances**, Board of Governors of the Federal Reserve System 2010,2013,2016. Available at <https://www.federalreserve.gov/econres/scfindex.htm> (accessed September 20, 2017).
- Vissing-Jørgensen, Annette**, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, 110 (4), 825–853.
- Yang, S. Alex and John R. Birge**, “Trade Credit, Risk Sharing, and Inventory Financing Portfolios,” *Management Science*, 2018, 64 (8), 3667–3689.
- Zinman, Jonathan**, “Household Debt: Facts, Puzzles, Theories, and Policies,” *Annual Review of Economics*, 2015, 7 (1), 251–276.

FIGURE I
OBSERVED CONSUMER GOODS INVENTORY

(a) Distribution of Inventory Levels




(b) Distribution of Inventory-to-Spending Ratio



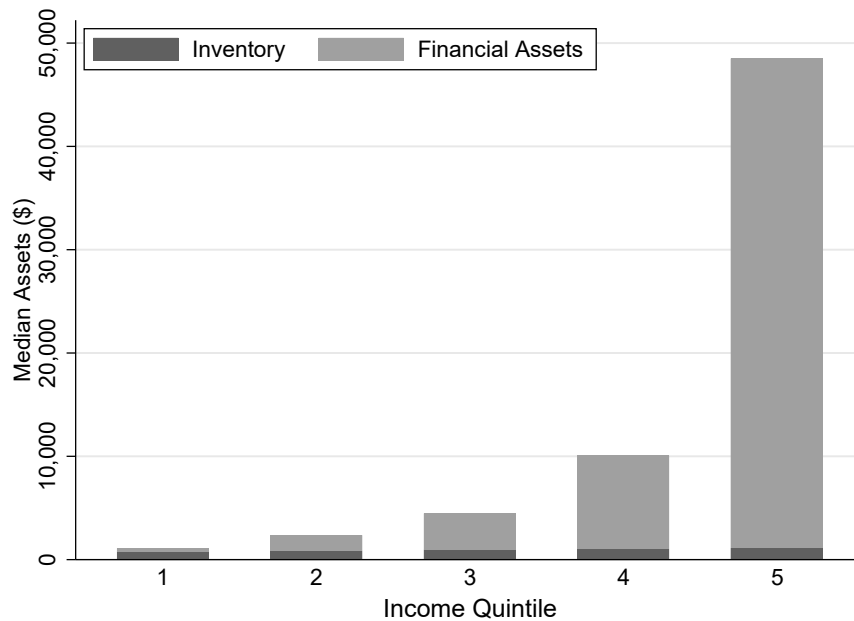
Notes: We compute the average value of inventory for households in the NCP over 2013 and 2014. Panel (a) plots how the average inventory level varies across households. Average inventory is plotted up to the 99th percentile. Summary statistics are computed using the raw data. Panel (b) plots the distribution of inventory as a share of the household's annual spending on goods covered by Nielsen. Both panels are constructed using Nielsen sampling weights.

FIGURE II
INVENTORY PORTFOLIO SHARE BY INCOME

(a) Inventory Portfolio Share

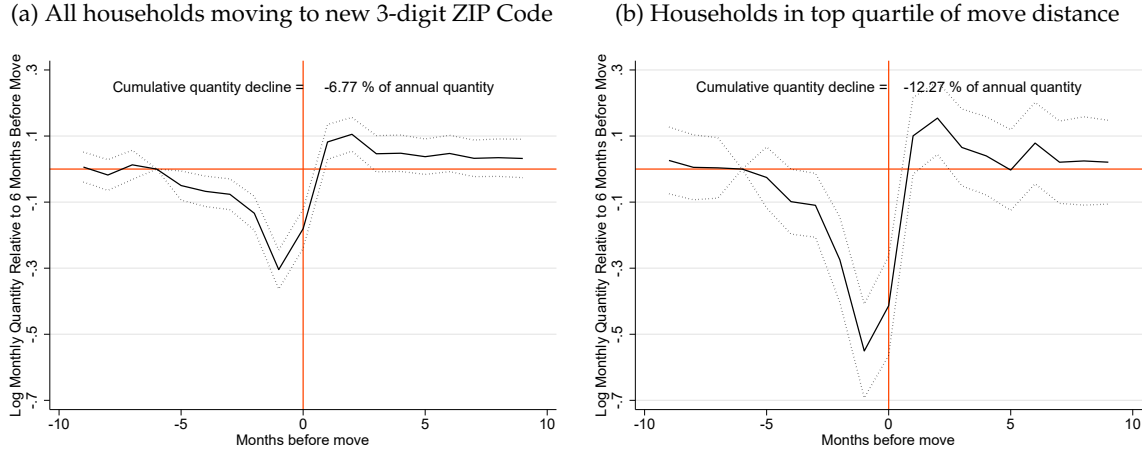


(b) Portfolio Composition



Notes: This figure is constructed by combining data from the NCP over 2013 and 2014 and the SCF over 2010, 2013 and 2016. We compute household income quintiles using the SCF and use household income reported by the Nielsen panelists to assign them to a quintile. We then compute the average value of inventory for each household and take the median across households in each quintile q (using Nielsen weights), $Inventory_q$. Finally, we compute the median level of financial assets held by the corresponding income quintile q in the SCF (using SCF weights) $Financial\ Assets_q$, and also the corresponding inventory portfolio share, $Inventory_q / (Financial\ Assets_q + Inventory_q)$.

FIGURE III
 VALIDATION: LOG QUANTITY PURCHASED AROUND MOVE DATES

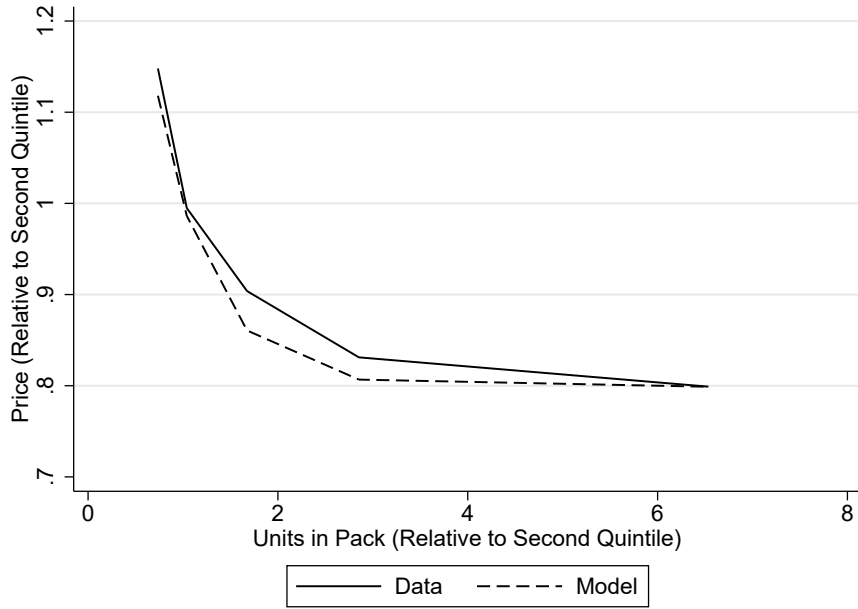


Notes: This figure shows the change in log quantity purchased around the time a household moves. For households who move to a new 3-digit ZIP Code in a given year we impute the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather than their old 3-digit ZIP Code). The figures plot estimates of β_s from the following specification and a 95 per cent confidence interval:

$$X_{i,t} = \sum_{s=-9, s \neq -6}^9 \beta_s \text{Moved}_{i,t-s} + \alpha \text{More than 9 months prior}_{i,t} + \gamma \text{More than 9 months after}_{i,t} + \text{Month FE} + \text{Household FE} + \epsilon_{i,t},$$

where X represents the log quantity purchased in ounces by household i in month t . $\text{Moved}_{i,t}$ is an indicator equal to 1 if household i moved in month t . $\text{More than 9 months prior}_{i,t}$ is an indicator equal to 1 if household i moved more than 9 months before month t , and $\text{More than 9 months after}_{i,t}$ is an indicator equal to 1 if household i moved more than 9 months after month t . The sample is restricted to 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 937km. Standard errors are clustered by household. Regressions are weighted using Nielsen sampling weights.

FIGURE IV
BULK CALIBRATION

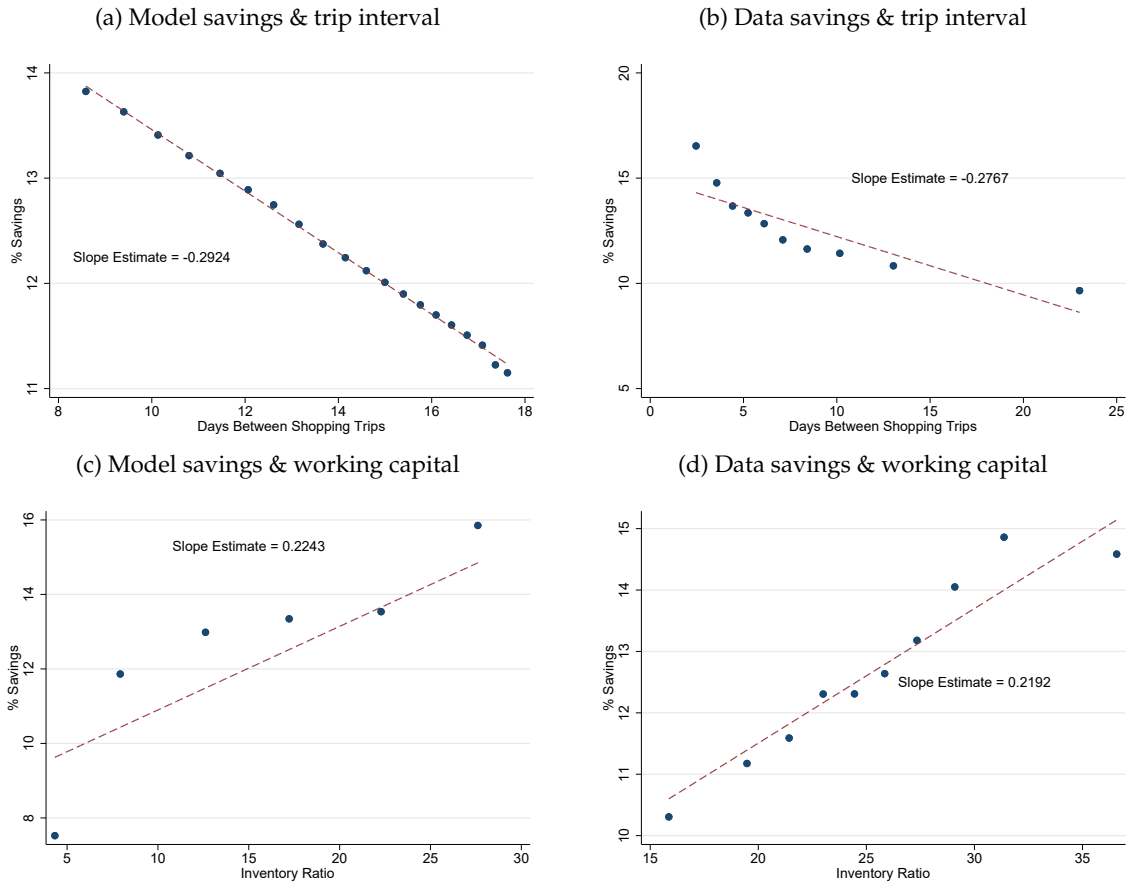


Notes: We first compute pack size quintiles for each product, and then compute the median number of units and unit price for each product and pack size quintile, weighted by UPC expenditure. We normalize both prices and units by dividing by the second quintile price and units. The range of available sizes varies substantially across products. As we ideally want to measure the aggregate obtained by increasing pack size uniformly across all products, we ensure that all products have a common range of normalized units. The 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 as 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 the following relationship:

$$\text{Price}_{p,g} = a_0 + a_1 e^{-\sigma \text{Units}_{p,g}} + \epsilon_{p,g}$$

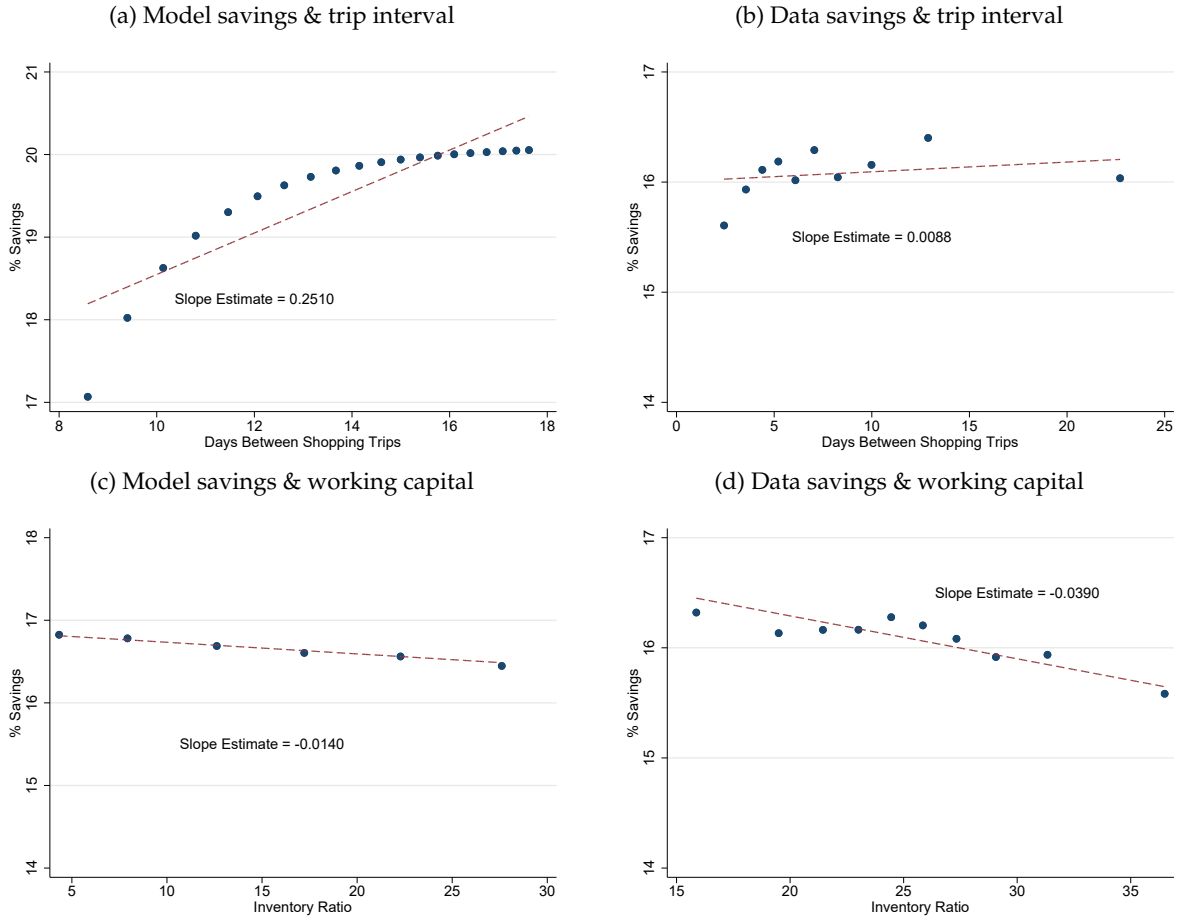
Where p is the product ID described in Section 4.8 and g is the pack size bin. We weight each observation by total spending on product p . The dashed line shows the relative price we assume in the model: $\text{Price} = \alpha + \beta e^{-\delta \text{Units}}$, where $\alpha = 0.80$, $\beta = 1.16$, $\sigma = 1.75$ and Units is the weighted average normalized units in units group g across all products. The solid line is constructed by computing the weighted average normalized retail price in units group g across all products.

FIGURE V
RELATIONSHIP BETWEEN DEAL SAVINGS, TRIP INTERVAL AND WORKING CAPITAL



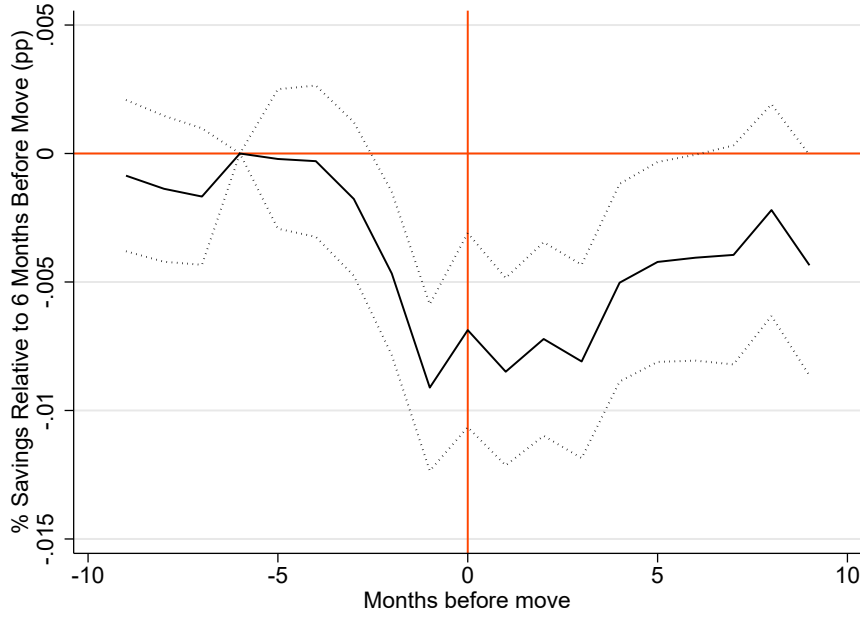
Notes: In Panel (a) we evaluate % deal savings and days between trips in the model for different values of k . Working capital is set sufficiently high that the working capital constraint does not bind. Days between trips is computed as $\Delta \times \frac{365}{12}$. In Panel (b) we plot the data relationship 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 data is the household's average inventory over a calendar year divided by annual Nielsen spending. Deal savings in the data are constructed using equation 25. In Panels (c) and (d) we control for the number of trips. Panels (b) and (d) are constructed using Nielsen sampling weights.

FIGURE VI
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 is 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 data is the household's average inventory over a calendar year divided by annual Nielsen spending. Bulk savings in the data are constructed using equation (26). 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 Nielsen sampling weights.

FIGURE VII
DEAL SAVINGS AROUND A MOVE



Notes: This figure shows the percentage point change in deal savings around the time a household moves. 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 β_s from the following specification and a 95 per cent confidence interval:


$$X_{h,i,t} = \sum_{s=-9, s \neq 6}^9 \beta_s \text{Moved}_{h,t-s} + \alpha \text{More than 9 months prior}_{h,t} + \gamma \text{More than 9 months after}_{h,t} \\ + \text{Month FE} + \text{Household} \times \text{Product Group FE} + \epsilon_{h,i,t},$$

where X is the percentage deal savings received on transaction i by household h in month t , winsorized at the first and 99th percentiles. We compute deal savings at the transaction level using the following formula:

$$X_{h,i,t} = \frac{\text{AvgPrice}_{z,u,y} - \frac{\text{Spending}_{h,i,t}}{Q_{h,i,t}}}{\text{AvgPrice}_{z,u,y}}$$

where z is the Zip Code in which the transaction occurred, u is the UPC purchased in the transaction and y is the calendar year of the transaction. $\text{Moved}_{h,t}$ is an indicator equal to 1 if household h moved in month t . $\text{More than 9 months prior}_{h,t}$ is an indicator equal to 1 if household h moved more than 9 months before month t , and $\text{More than 9 months after}_{h,t}$ is an indicator equal to 1 if household h moved more than 9 months after month t . The sample is restricted to households who moved to a new 3-digit ZIP Code exactly once. 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 Nielsen sampling weights.

FIGURE VIII
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. Each point 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 dotted line shows predicted values from $\frac{\text{Annual Savings}}{\text{Annual Spending}}_h = \alpha + \beta \frac{\text{Annual Avg. Inventory}}{\text{Annual Spending}}_h + \gamma \text{Trips}_h + \epsilon_h$. The regression is weighted using Nielsen sampling weights.

TABLE I
VALIDATION: RELATIONSHIP BETWEEN DURABILITY AND INVENTORY RATIO

	(1)	(2)	(3)	(4)
Semi-perishable	4.8*** (0.0)	5.2*** (0.0)	4.8*** (0.0)	5.0*** (0.0)
Non-perishable	8.5*** (0.0)	8.7*** (0.0)	8.6*** (0.0)	8.6*** (0.0)
Potential Bulk Savings			6.7*** (0.1)	6.5*** (0.1)
Number of Trips (100s)			-7.7*** (0.1)	
Household FE		X		X
Number of Observations	1,792,273	1,792,273	1,675,238	1,675,238

Notes: This table combines data from the NCP over 2013 and 2014; the SCF over 2010, 2013 and 2016; and the FSIS. We estimate the following regression specification, where h indexes households and c indexes product perishability categories (perishable, semi-perishable and non-perishable):

$$\text{Inventory Ratio}_{h,c} = \beta_1 \text{Semi-perishable}_c + \beta_2 \text{Non-perishable}_c + \beta_3 X_{h,c} + \text{Household FE} + \epsilon_{h,c}. \quad (29)$$

$\text{Inventory Ratio}_{h,c}$ is the ratio of household inventory to annual spending in perishability category c . Columns 1 and 3 show results without household fixed effects. The base perishability category is “Perishable”. These are products which have a life of less than three months. We define “Semi-perishable” items as those with a life of more than three months and less than one year. Coefficients are multiplied by 100. Standard errors are clustered by household. Regressions are weighted using Nielsen sampling weights. * $p < .1$, ** $p < .05$, *** $p < .01$.

TABLE II
MODEL CALIBRATION

Name	Parameter	Calibrated value	Source/target
Trip fixed cost	k	\$4.85	Baker et al. (forthcoming) .
Deal probability	x	0.239	x is calibrated to match NCP deal share.
Full list price	p_f	1.065	p_f and p_d are jointly calibrated to match average discount size in the NCP and $E[p] = xp_d + (1 - x)p_f = 1$.
Deal price	p_d	0.794	
1 - max bulk savings %	α	0.799	α , β and σ are jointly calibrated to match the relationship between pack size and unit price in the NCP.
Bulk savings parameter	β	1.157	
Bulk savings parameter	σ	1.75	
Monthly consumption	C	492.20	C and δ_0 are jointly calibrated to match average annual spending and inventory in the NCP.
Storable depreciation	δ_0	0.004	
Perishable depreciation	δ_1	2.885	Sets expected expiry to average expiration date for products lasting < 1 month (0.347 months)
Non-perishable share	good s_0	0.769	NCP expenditure share of goods with expiration date ≥ 1 month

TABLE III
FINANCIAL RETURNS TO HOUSEHOLD INVENTORY INVESTMENT

Working Capital	Inventory Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings:		Interval Δ^* (Months)	m_0^*	Net Return (%)
				Deal	Bulk			
0.05	0.04	136.90	101.04	9.54	14.71	0.24	4	82.18
0.10	0.09	363.91	113.93	11.99	17.03	0.28	7	22.92
0.15	0.14	601.66	113.47	13.14	17.11	0.28	10	7.37
0.20	0.19	844.26	112.23	13.66	17.03	0.28	13	1.27
0.25	0.23	1034.41	114.64	13.79	17.28	0.29	15	0.00

Notes: This table is constructed by solving the model for different values of \bar{I} . The working capital ratio in column 1 is the ratio of working capital in \$ (\bar{I}) to unconstrained annual model spending ($\$399.62 \times 12 = \$4,795$). The inventory ratio is the ratio of average inventory to annual model spending. Maximum cash is the cash held immediately prior to a shopping trip. This is also the same as the total price paid in store in \$, $\sum_l P_l \cdot S_l$. Immediately following a trip, cash in the model is equal to zero (because there are no other motives to hold cash, such as precautionary or speculative motives). Average cash holding in the model is therefore half of the maximum holding, and the maximum cash holding is equal spending per trip. Minimum inventory is the inventory level immediately prior to a shopping trip. The maximum level of inventory is equal to the sum of minimum inventory and maximum cash and is reached immediately after a shopping trip. Deal savings are % savings of annual spending due to buying an item on sale. Bulk savings are % savings of annual spending due to buying a larger pack size. Interval Δ^* is the optimal length of time between trips measured in months. m_0^* is the optimal deal shopping strategy for non-perishable goods, subject to the household working capital constraint. The net return incorporates not only in-store savings but also depreciation and trip fixed costs.

TABLE IV
FINANCIAL RETURNS WITH HIGH SHOPPING TRIP FIXED COST ($k = \$100$)

Working Capital	Inventory Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings:		Interval Δ^* (Months)	m_0^*	Net Return (%)
				Deal	Bulk			
0.05	0.02	31.56	207.10	3.07	19.76	0.44	1	154.14
0.10	0.07	265.52	212.28	8.37	19.87	0.47	4	48.09
0.15	0.11	499.79	217.01	10.05	19.91	0.48	6	20.91
0.20	0.14	674.26	239.84	10.35	19.99	0.52	7	14.43
0.25	0.19	956.21	237.04	11.12	19.98	0.52	9	4.26
0.30	0.24	1210.93	224.91	11.71	19.95	0.50	11	3.06

Notes: See the description in Table III.

Financial Returns to Household Inventory Management

Scott R. Baker

Stephanie Johnson

Lorenz Kueng

A Constant Consumption Assumption

In order to compute household inventories we make the assumption that each household has constant consumption at the product group level. In this section, we show how violations of this assumption influence the inventory calculation. Although non-constant consumption does lead to inventory being overstated, assuming several plausible non-constant consumption patterns we show that the effect is small relative to the overall level of inventory we find.

There are two main ways in which we expect the constant consumption assumption to be violated. First, households may have non-constant aggregate grocery consumption, for example due to holidays or parties. This violation would not be addressed by aggregating across products.

The second type of violation occurs when households do not have constant consumption at the assumed level of aggregation (holding their aggregate consumption constant). For example, if consumption is assumed to be constant at the UPC level, but households regularly switch brands, pack sizes or substitute very similar products from week to week. Seasonal consumption of certain products (such as turkey or stuffing mix) also falls in this category. The question is then which level of aggregation is appropriate, and, at the chosen level of aggregation, what is the likely degree of inventory overstatement.

Below we provide some examples illustrating how violations of these assumptions affect the inventory calculation. In all examples we assume that true consumption is equal to spending and true inventory is therefore zero. We then compute the ratio of measured inventory to annual spending under the incorrect assumption that consumption is constant. Figure [B.4a](#), shows a household with a large spike in consumption at four dates spread throughout the year. Consumption on these ‘celebration’ days is five times consumption on a typical day. This pattern of non-constant consumption yields a computed ratio of inventory to annual spending equal to 0.007. Figure [B.4b](#) shows consumption for a household who consumes around 25% more in the months of June and July than it does at other times. Annual spending is the same as in Figure [B.4a](#). This consumption pattern yields a ratio of inventory to annual spending equal to 0.019.

Next we consider the case where aggregate consumption is constant, but households switch between product groups. Consequently there are large fluctuations in consumption at the product group level. In Figure [B.5a](#), households consume two product groups and alternate between them each day. This pattern yields an inventory ratio of 0.004 if consumption is assumed to be constant. In Figure [B.5b](#), households alternate between product groups each week. This yields an inventory ratio of 0.012.

These examples illustrate that the inventory calculation is generally robust to fairly extreme violations of constant consumption, such as occasional large parties, very seasonal consumption, and

extreme switching between product groups for variety on a day-to-day or week-to-week basis.

While it is challenging to provide direct empirical evidence on households' actual consumption of the products we consider, we use data from NHANES which provide information on food and beverage items consumed by an individual on non-consecutive days. Figure B.6a shows the average share by number of days between interviews without any aggregation (that is, the share of Day 1 NHANES food items which were also consumed on Day 2).

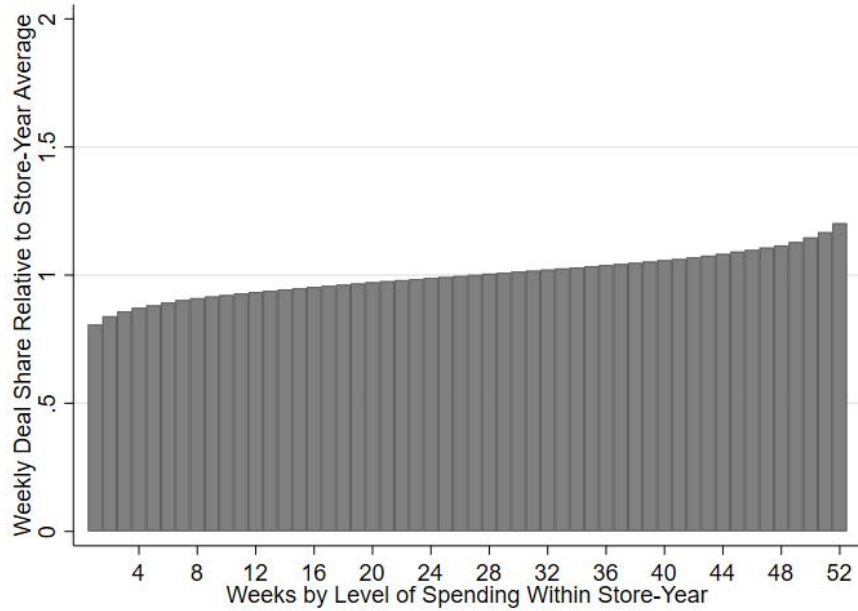
We also separate respondents where one interview day was a weekday and the other day was a weekend. Even without further aggregation the share of items also consumed on Day 2 is quite high, at around 40%. The share declines slightly with the time between interviews, consistent with some of the persistence being driven by households consuming items from the same shopping trip, but remains high even with a gap of 9 days. Figure B.6c shows the effect of aggregating to Nielsen product group. In this case around 60% of Day 1 product groups were also consumed on Day 2. After aggregating to Nielsen department (Figure B.6e) this further increases to 90%.

NHANES respondents also report the amount of each item consumed in grams. Figures B.6b, B.6d and B.6f illustrate the relationship between Day 1 quantity and Day 2 quantity for the same item, product group or department. Regardless of the level of aggregation, Day 1 quantity is closely related to Day 2 quantity (the coefficient is also close to one when excluding items where a very large amount is consumed in Day 1).


B Additional Appendix Figures and Tables

FIGURE B.1
RETAILER DEAL CONCENTRATION

(a) Ranked Weeks



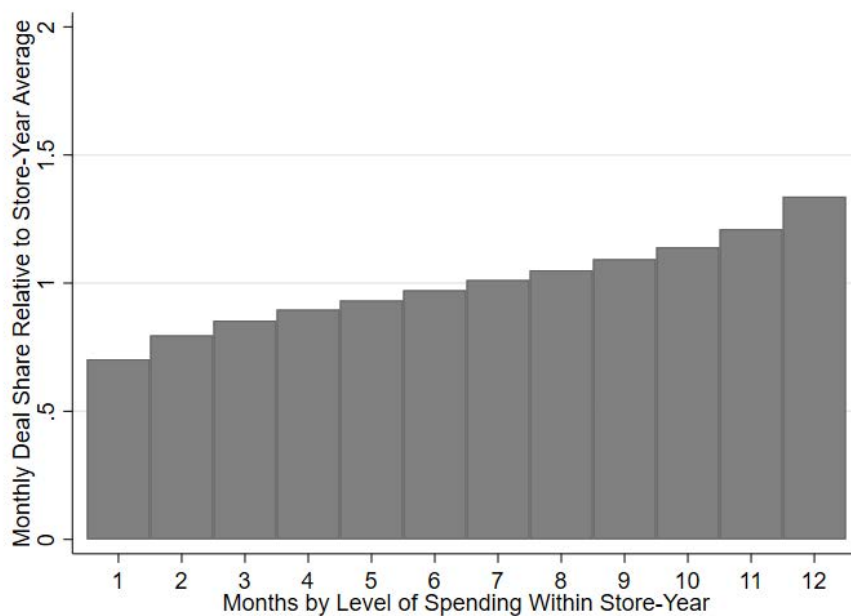
(b) Calendar Weeks




Notes: We compute the share of deal sales for each retailer in each week using the deal flag in the NCP (which includes both coupon and non-coupon deals), and then divide by the retailer's average deal share over the year. Panel (a) plots the average across retailers by ranked weeks (so week 1 is the week with the lowest deal share). Panel (b) plots the average by calendar week. We restrict the sample to large retailers with more than 1000 separate items sold each week to NCP households.

FIGURE B.2
STORE DEAL CONCENTRATION

(a) Ranked Months

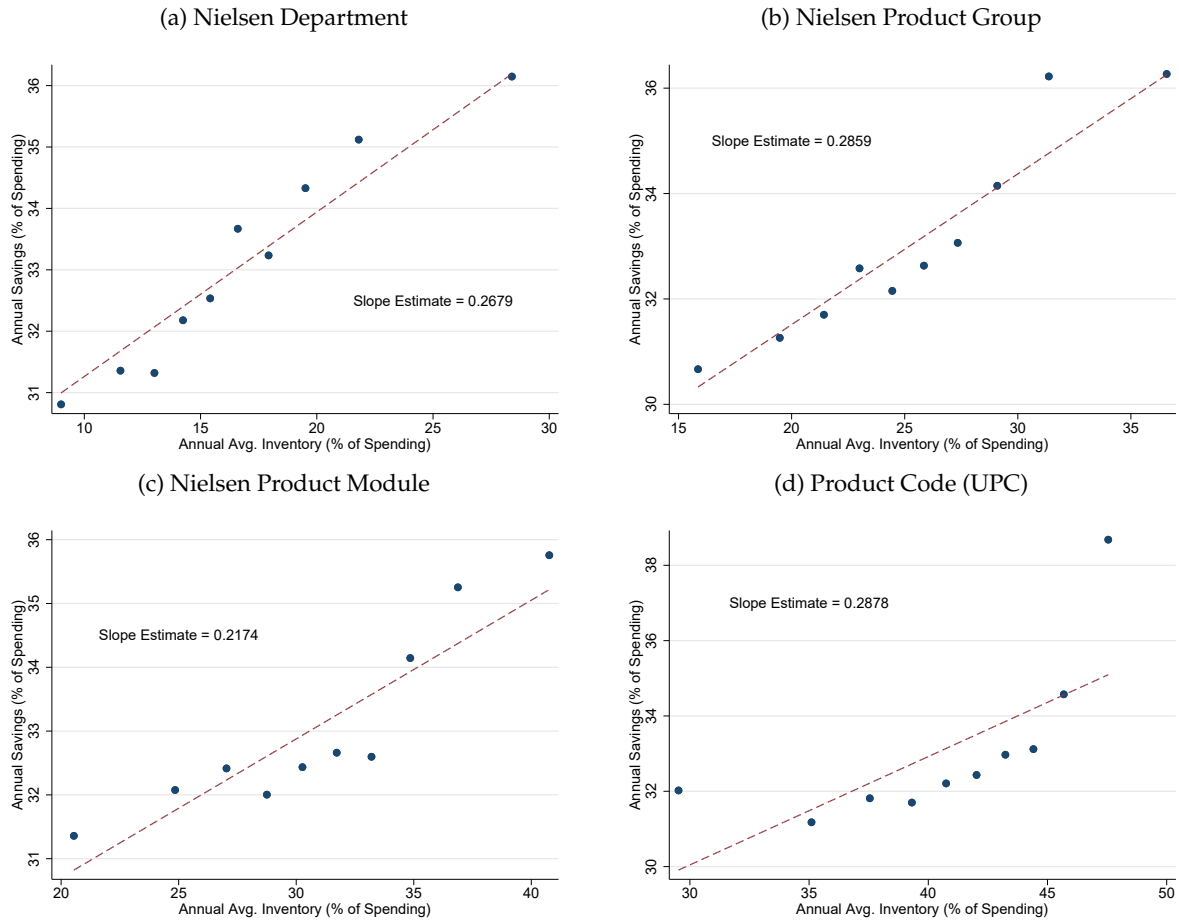


(b) Calendar Months



Notes: We compute the share of deal sales for each store in each month using the deal flag in the NCP (which includes both coupon and non-coupon deals), and then divide by the store's average deal share over the year. Panel (a) plots the average across stores by ranked month (so month 1 is the month with the lowest deal share). Panel (b) plots the average by calendar month. We restrict the sample to large stores with more than 500 separate items sold each month to NCP households.

FIGURE B.3
GROSS RETURN UNDER ALTERNATIVE AGGREGATION ASSUMPTIONS



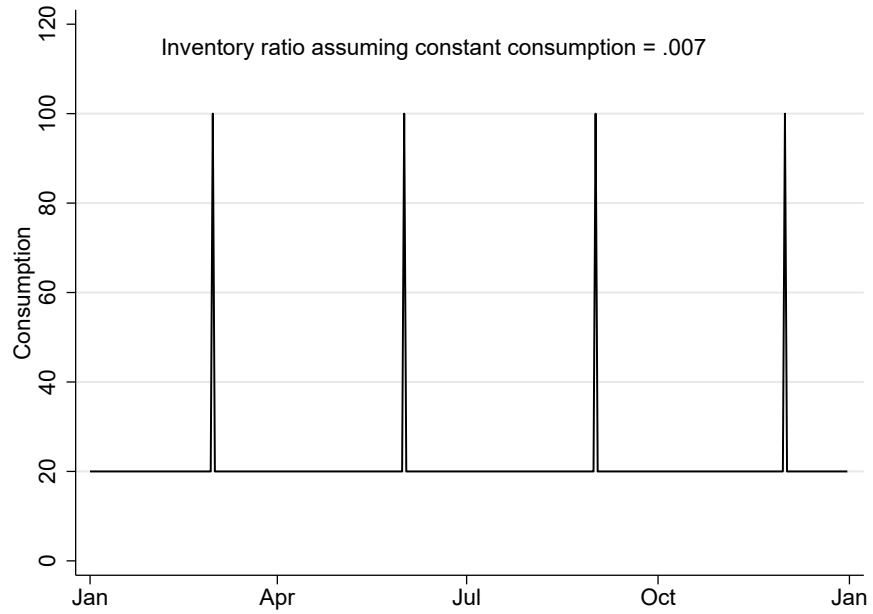
Notes: This figure uses the NCP over 2013 and 2014 to illustrate the relationship between in-store savings and average inventory as a percentage of spending. Average inventory is computed with decreasing level of aggregation, starting from the 10 “Nielsen Departments” and going to 118 “Nielsen Product Groups” (our main specification in the text), 1,305 “Nielsen Product Modules” and finally to the extreme of using over 2 million unique product codes (UPC). Each point 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 dotted line shows predicted values from the following regression specification:

$$\frac{\text{Annual Savings}}{\text{Annual Spending}}_h = \alpha + \beta \frac{\text{Annual Avg. Inventory}}{\text{Annual Spending}}_h + \gamma \text{Trips}_h + \epsilon_h.$$


Regressions are weighted using Nielsen sampling weights.

FIGURE B.4
NON-CONSTANT AGGREGATE CONSUMPTION

(a) Example 1



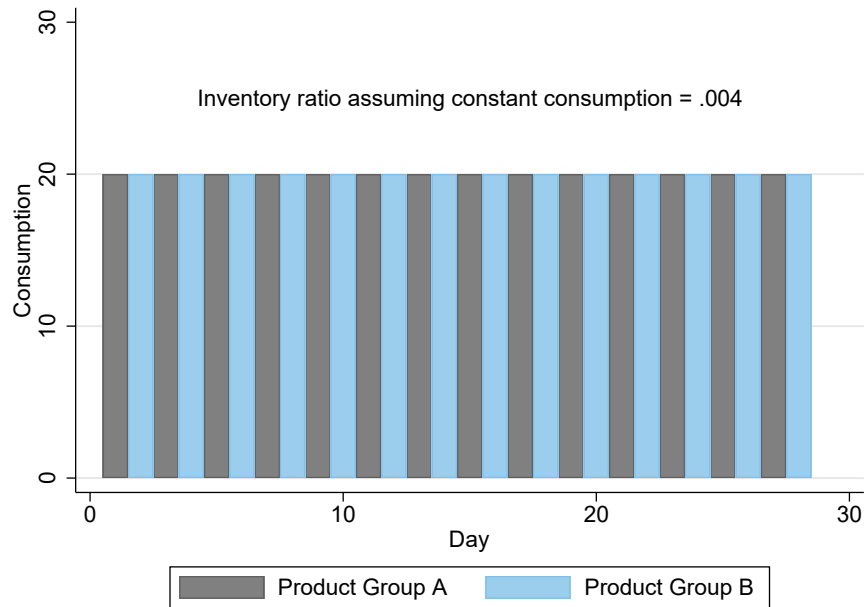
(b) Example 2




Notes: Panel (a) plots consumption for a household who consumes five times more on the first day of March, June, September and December than on other days of the year. Panel (b) plots consumption for a household who consumes around 25 per cent more in June and July than in other months. The inventory ratio assuming constant consumption is computed using the method described in Section 3.

FIGURE B.5
PRODUCT GROUP SWITCHING

(a) Example 1

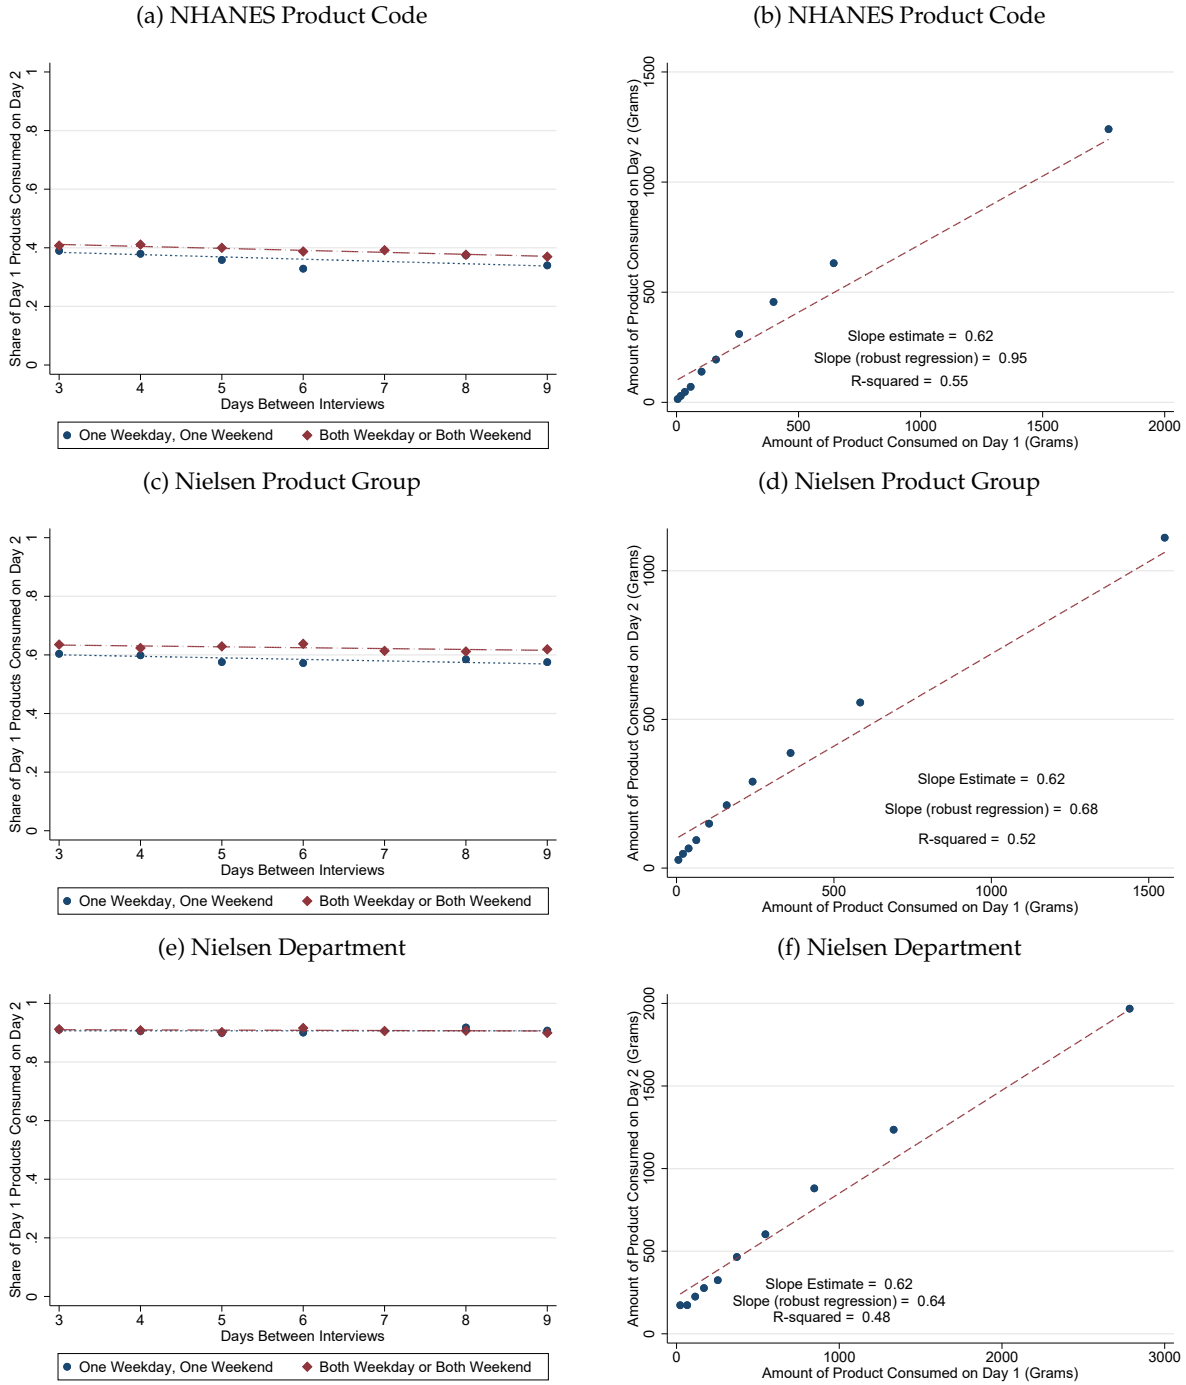


(b) Example 2



Notes: Panel (a) shows the consumption pattern of a household who consumes only Product Group A on one day and only Product Group B on the following day. Panel (b) shows the consumption pattern of a household who consumes only Product Group A one week and only Product Group B the following week. Aggregate consumption is constant. The inventory ratio assuming constant consumption at the product-group level is computed using the method described in Section 3.

FIGURE B.6
 VALIDATION: CONSUMPTION PERSISTENCE



Notes: For each individual we enumerate the NHANES products, Nielsen product groups, and Nielsen departments consumed on Day 1 and Day 2 of the survey. We then compute the share of Day 1 products, groups or departments which were also consumed on Day 2. Panels (a), (c) and (e) show the average share by number of days between interviews. Because consumption patterns may differ on weekdays and weekends, we also show results separately for individuals where one survey day was a weekday and the other was on the weekend. Panels (b), (d) and (f) show the relationship between the amount of an NHANES product, Nielsen product group, or Nielsen department consumed on Day 1 and the amount consumed on Day 2.

TABLE B.1
 RETURNS UNDER ALTERNATIVE AGGREGATION ASSUMPTION: NIELSEN DEPARTMENT

Working Capital	Inventory Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings:		Interval Δ^* (Months)	m_0^*	Net Return (%)
				Deal	Bulk			
0.05	0.04	134.58	98.46	9.53	15.33	0.24	4	67.78
0.10	0.09	356.42	110.82	11.99	17.34	0.28	7	14.63
0.15	0.14	582.87	109.20	13.17	17.26	0.28	10	0.98
0.20	0.15	662.89	109.15	13.39	17.26	0.28	11	0.00

Notes: This table shows results when the model is calibrated to match average inventory of \$716 rather than \$1,113. This is the average inventory obtained when we use Nielsen “departments” rather than “product groups” to compute inventories. See the description in Table III.

TABLE B.2
RETURNS UNDER ALTERNATIVE AGGREGATION ASSUMPTION: NIELSEN PRODUCT MODULE

Working Capital	Inventory Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings:		Interval Δ^* (Months)	m_0^*	Net Return (%)
				Deal	Bulk			
0.05	0.04	135.62	100.76	9.54	14.22	0.24	4	89.31
0.10	0.09	359.95	113.25	11.99	16.70	0.28	7	26.24
0.15	0.14	598.60	113.44	13.12	16.88	0.28	10	9.54
0.20	0.19	833.17	111.26	13.66	16.70	0.28	13	4.03
0.25	0.23	1041.26	115.96	13.75	17.22	0.29	15	1.30
0.30	0.29	1302.80	115.89	13.88	17.22	0.29	18	0.00

Notes: This table shows results when the model is calibrated to match average inventory of \$1,379 rather than \$1,113. This is the average inventory obtained when we use Nielsen “product modules” rather than “product groups” to compute inventories. See the description in Table III.

TABLE B.3
RETURNS UNDER ALTERNATIVE AGGREGATION ASSUMPTION: UPC

Working Capital	Inventory Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings:		Interval Δ^* (Months)	m_0^*	Net Return (%)
				Deal	Bulk			
0.05	0.03	74.63	89.14	7.77	15.32	0.30	3	87.33
0.10	0.08	230.86	94.36	11.01	16.74	0.33	6	34.45
0.15	0.13	399.55	90.46	12.55	16.40	0.33	9	13.64
0.20	0.19	568.56	86.97	13.28	15.94	0.32	12	6.97
0.25	0.23	727.47	91.27	13.38	16.66	0.33	14	2.89
0.30	0.28	859.19	91.14	13.51	16.66	0.33	16	2.18
0.35	0.33	1051.89	90.46	13.62	16.57	0.33	19	0.71
0.40	0.38	1191.42	91.01	13.63	16.66	0.33	21	0.49
0.45	0.44	1383.27	90.41	13.67	16.57	0.33	24	0.15
0.50	0.48	1525.19	90.98	13.67	16.66	0.33	26	0.06
0.55	0.52	1658.83	90.97	13.67	16.66	0.33	28	0.01
0.60	0.56	1792.52	90.97	13.67	16.66	0.33	30	0.00

Notes: This table shows results when the model is calibrated to match average inventory of \$1,863 rather than \$1,113. This is the average inventory obtained when we use UPCs rather than “product groups” to compute inventories. See the description in Table III.