FINANCIAL RETURNS TO HOUSEHOLD INVENTORY MANAGEMENT *

Scott R. Baker **Stephanie Johnson** **Lorenz Kueng**

July 2022

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 $900 on average. Such holdings can eclipse total financial assets among households in the lowest income quintile. 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. Inventory management of consumer goods provides one alternative to investments in risky financial markets at low levels of liquid wealth.

JEL Classification: G51, G11, D14, D13, D12, D11, E21

Keywords: household working capital, financial returns, inventory, stockpiling

*Baker: Northwestern University, Department of Finance, 2169 Campus Drive, Evanston, IL 60208; s-baker@kellogg.northwestern.edu. Johnson (corresponding author): Rice University, Jones Graduate School of Business, Janice and Robert McNair Hall, 1900 Rice Boulevard, Houston, TX 77005; stephanie.g.johnson@rice.edu. Kueng: Università della Svizzera italiana (USI Lugano), Swiss Finance Instutute (SFI), and CEPR, Via Giuseppe Buffi 13, CH-6900 Lugano TI, Switzerland; lorenz.kueng@usi.ch. We thank Alex Butler, Hector Calvo-Pardo, James Choi, Giuseppe Di Giacomo, John Gathergood, Gustavo Grullon, Luigi Guiso, Michael Haliassos, Xavier Jaravel, Pierre Mabille, Soren Leth-Petersen, Jonathan Parker, Luigi Pistaferri, Tarun Ramadorai, Antoinette Schoar, and Qingyuan Yang for valuable feedback and our discussants Francesco d'Acunto, Andres Almazan and Matteo Benetton and participants at seminars at Einaudi Institute for Economics and Finance (EIEF), National University of Singapore (NUS), Northwestern University, Rice University, CEPR European Conference on Household Finance 2020, Copenhagen New Consumption Data Conference, Western Finance Association Meeting 2021, 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. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

While a large number of American households hold small amounts or even zero financial assets, all households hold at least some resources in the form of consumer goods inventories. These inventories can be managed over time through strategic shopping behavior as households are able to take advantage of coupons, temporary low prices at retailers, and savings from buying in bulk.

In this paper, we study how the financial return to investment in inventories affects households' desire to hold liquid assets like cash and cash equivalent assets (e.g. checking accounts, transaction accounts, credit card lines of credit). We refer to these combined resources – the sum of cash and inventory – as *household working capital*.[1]

The paper makes two main contributions. First, using scanner data from AC Nielsen and income and asset data from the Survey of Consumer Finances (SCF), we provide estimates of this hitherto neglected source of non-financial wealth on households' balance sheets. Aggregating across all Nielsen Homescan goods, we find that households hold on average nearly $1,000 in consumer goods inventory at any given time. For households in the lowest quintile of annual household income (under $22,000), inventory likely represents a greater store of value than total household financial assets.

Second, we build a parsimonious model of inventory management to compute the marginal financial return to investing in household working capital through the maintenance of liquid savings and engaging in strategic shopping behavior. These returns net of product deprecation and trip costs are approximately risk-free and well above 20% at low levels of working capital, though they decline rapidly as inventory levels increase.

Optimal inventory management therefore 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. For instance, Orhun and Palazzolo (2019) note that low-income households are less responsive than higher income households to sales or promotions in part due to a lack of liquidity reserves to employ for intertemporal substitution. This stockpiling behavior of storable non-durables can produce can drive volatile expenditures alongside smooth consumption paths, mirroring a tendency recognized among durable goods (see Parker (1999), Browning and Crossley (2001) and Browning and Crossley (2009)). The high returns observed in our data can, in a minority of instances, even rationalize high cost borrowing like credit card debt.

When computing 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. We also extend our framework to include the costs from product depreciation and spoilage and the relation between the level of inventory holdings and differences in shopping trip fixed costs associated with different shopping behaviors. Moreover, relative to the previous literature, we endogenize the timing of shopping trips.

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

[1] This combination of financial resources and consumer goods echoes firms' working capital which includes both current account resources as well as materials and inventories that may be at least partially non-tradeable.

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 in the model optimally chooses shopping trip frequency to minimize the cost of providing a given consumption stream, subject to a household working capital constraint. The model therefore allows us to study how "investing" in household working capital generates a return in the form of reduced trip costs and lower per unit prices, taking into account product depreciation.

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 is on sale at any given time across a household's total basket. This formulation is derived from an assumption of independent price deals across goods and backed by observations from the data. It has implications for households' cash and cash-equivalent 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.[2]

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 the stockpiling channel of consumer responses to these shocks. In Baker, Kueng, McGranahan and Melzer (2019) and Baker, Johnson and Kueng (2021), we use anticipated local sales tax increases at a monthly frequency and show that consumers respond strongly along several margins, including stocking up on products while taxes are still low. Hendel and Nevo (2006) identifies the related effect of temporary price discounts on household inventory in the context of a dynamic discrete choice model of individual product demand with exogenously given shopping trip frequency. The authors highlight the importance of household inventory to explain the large price elasticities observed in scanner data.

By highlighting the role of household working capital for households' portfolio allocation and spending behavior – 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); Rampini (2019)), inventories of consumer goods and household working capital has been largely ignored by the household finance literature.[3]

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

[2] In contrast, if aggregated deals are autocorrelated, households may want to hold substantial additional cash or cash-equivalents 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 C.1 and C.2.

[3] A 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.

Bertaut and Starr (2000), who study U.S. households' portfolios.[4] 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.[5] We also note that adding household inventory management to a household's portfolio choice problem can potentially affect its decision of whether to participate in financial markets. Household working capital could therefore provide another partial explanation to the stock market participation 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)).[6]

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

Finally, 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 rationalize some of the borrowing at fairly high interest rates that lower income households engage in (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)).

[4]The 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.)."

[5]There is a literature in macroeconomics studying heterogeneity in the effective price paid for similar goods across households and over 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).

[6]The literature on the participation puzzle is among the oldest in household finance and too large to adequately survey here. The handbook chapters by Guiso and Sodini (2013) and Beshears, Choi, Laibson and Madrian (2018) provide recent surveys. A non-exhaustive list of explanations of the participation puzzle include pecuniary and non-pecuniary participation fixed costs (Luttmer (1999); Vissing-Jørgensen (2002)); low financial literacy (Van Rooij, 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 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)).

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. 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. They utilize regular reminders to participants that prompt them to report fully and will remove non-compliant households from their panel. Broda and Weinstein (2010) and Einav, Leibtag and Nevo (2010) perform more analysis of the NCP. Overall, they deem the NCP to be of comparable quality to many other commonly-used self-reported consumer data. The NCP primarily covers trips to grocery, pharmacy, and mass merchandise stores but also spans a wider range of channels such as catalog and online purchases, liquor stores, delis, and video stores. 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. To the extent we do not observe household purchase or stockpiling of clothing, electronics, or other larger purchases, we will underestimate inventory and thus consider our estimates a lower bound of household inventory.

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. This data covers almost 100 retail chains with over 40,000 unique stores in over 350 MSAs across the country.

In general, the data span many of the largest retailers in the grocery, mass merchandiser, drug-store, and pharmacy 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 retailers.

In total, these data comprise over 10 billion transactions per year worth nearly $250 billion.

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 2013-2014 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. The item information is very detailed, for example distinguishing between whether a vegetable item was canned, fresh or frozen. In some cases the NHANES code corresponds to a meal with multiple ingredients (eg. "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.

3 Measuring Levels of Household Inventory

To compute household inventories using the NCP data we must make some additional assumptions. Although we can track the flow of purchases for individual products over time in the NCP – including the exact time stamp of each purchase – we must make three assumptions which are necessary because both the initial inventory and the flow of consumption are unobservable.

Our first assumption is that a household has just enough in stock initially to ensure that inventory does not become negative at any point in the sample period. This implies that inventory hits zero at least once. The second assumption is that a given product's rate of inventory depletion (rate of consumption and depreciation) is constant throughout the year. Using direct evidence from NHANES on actual consumption rather than spending, Appendix B shows that consumption is indeed fairly constant when aggregating up to Nielsen product groups (see also Aguiar and Hurst (2013) for similar evidence). The third assumption is that average annual inventory depletion equals average annual spending such that inventory is stationary. With these three assumptions, we can compute the initial

level of inventory for each household in the NCP sample as well as the inventory level at all later points in time during which we observe the household.

How well these assumptions recover unobserved true household inventories depends on the level 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: We most definitely overestimate inventory levels if we do not aggregate individual products, and we likely underestimate inventory if we aggregate them too coarsely. Next, we briefly provide the intuition and motivate our preferred product aggregation.

On the one hand, if we do not aggregate individual products (UPCs) at all, then our assumption that consumption is constant throughout the year is poor and inventory is overstated. For instance, if a household switches cereal products sequentially to try more varieties (say between Kellogg's Raisin Bran Original and Kellogg's Raisin Bran Crunch), they do not hold a stockpile of all varieties at once, but rather consume them one after the other. If the household consumes the same amount of cereal every day, consumption is constant at the product *group* level (cereal) but not at the UPC level.

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

For these reasons, the average level of measured inventory varies with the level of product aggregation we assume. As outlined in Section 2.2, Nielsen provides several product classifications which serve as potential levels of aggregation. We impose our three assumptions for each individual product category and sum up inventories over all categories to get total household inventory.

With this approach, the average amount of household inventory in NCP goods aggregated to "product groups" is \$867. This is our preferred level of aggregation, which still leaves us with fairly disaggregated data.[7] Aggregating further to the 10 "Nielsen Departments" would yield an average inventory of \$536, which very likely underestimates inventory for the reasons outlined above. Aggregating less by using Nielsen's 1,305 "Product Modules" or even all approximately 2 million individual UPCs would yield average inventories of \$1,147 and \$1,650, respectively. In our opinion, these latter two are likely to overstate inventories as the constant consumption assumption is probably inappropriate at such aggregation levels, especially for individual UPCs. Unless stated otherwise, in the remainder of the paper we therefore aggregate UPCs to "product groups", and we document how our main findings change for different levels of product aggregation in Appendix C.

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

[7] This 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".

3.1 Computing Inventories

We derive a formula for average household inventory using a value-based approach. This uses dollars of spending to measure the inflow of inventory. However, under this approach, fluctuations in unit prices over time may in principle affect the accuracy inventory calculation. An alternative approach would be to derive average inventory in terms of quantity, and then apply an average annual per unit price to convert obtain inventory in dollars. Given that in practice both approaches yield very similar results on a consistent set of UPCs, we prefer the value approach as it allows us to obtain 100% coverage of NCP products, including those where the unit of measurement is not directly comparable across items.

The average inventory held over the period from time zero to T is $\bar{I}_T = \frac{1}{T}\int_0^T I(t)dt$, where $I(t)$ is the unobserved level of inventory at time t. Inventory at time t reflects the initial time zero level of inventory $I(0)$, purchases made on trips between time 0 and time t, and the rate of inventory depletion d, which we assume to be constant:[8]

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

where $\{t_j\}$ are the dates of the n_t shopping trips occurring between time 0 and time t (i.e. time stamps in the NCP data). S_{t_j} is the value of purchases made on the jth trip. $I(0)$ contributes to average inventory from time 0 to T and each trip size S_{t_j} contributes from time t_j to T so that the integral is:

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

Hence, average inventory is:

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

When applying the formula to each product group g for each household h in the data, we compute average annual inventory. Hence, with time measured in years, we have $T = 1$ and the time stamp t_j of trip j is relative to the start of the year and takes values between 0 and 1. Assuming annual depletion is equal to annual spending, annual average inventory of household h in product group g in calendar year y is:

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

where $n_{y,h}$ is the number of trips household h makes over calendar year y and $1 - t_j$ is the share of the calendar year remaining when trip j occurs.

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

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

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

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

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

Note that $I(0)_{y,h,g} \geq 0$, because $I(t_1)_{y,h,g} \leq 0$.

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 and furniture and also excludes most clothing and electronics. On the household balance sheet, some such durable items would be classified as long-term physical assets – corresponding to "Property, Plant, and Equipment (PP&E)" on the corporate balance sheet – and are therefore not included in our definition of household working capital. Figure Ia shows the distribution in inventory in dollars. Figure Ib shows the distribution of the 'inventory ratio': the value of inventory divided by annual household spending on the same set of Nielsen products. The median household holds inventory worth 0.2 years of spending (2.4 months).

Note that 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 the model of Section 4, we will use FSIS data to calibrate the product depreciation rates in our model, which allows us to explicitly take spoilage into account when computing financial returns.

3.2 Inventory Portfolio Shares

Overall, inventory is an important asset for many households, even with durables excluded. 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 63% of financial 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 can drive a non-trivial portion of observed spending fluctuations within a year.

3.3 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 that we obtain from FSIS, a third-party source. Second, we show that for households run down inventories in advance of a move. 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.

3.3.1 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 group a usable life in months (between 0.13 and 60) using FSIS data. The majority of spending in the NCP is on products with a lifetime of more than 6 months (around 70%).

The relationship between inventory and shelf-life is increasing and concave. The relationship is also robust to controlling for the household's trip frequency, which also has substantial explanatory power for the inventory ratio. Starting from a shelf-life of zero, a one month increase in shelf-life raises the inventory ratio by around 0.4 percentage points. Column 5 shows an alternative specification with an indicator equal to one for products with a shelf-life greater than 6 months. Households hold around 6 percentage points of annual spending more inventory of these products.

3.3.2 Inventory Dynamics of Movers

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

Figure III shows that the behavior of the subset of households who move is consistent with this conjecture. Households cut purchase quantities well in advance of a move and by a large amount, consistent with our finding that they hold a large stockpile of inventories.[9] Purchase quantities return to normal immediately following the move. An alternative explanation for the decline in purchases prior to a move is that households are cutting consumption to pay for move-related expenses – for example to cover transportation costs, down payments, or security deposits. However, we also find that realized in-store savings decline around the move, which seems inconsistent with this story but is consistent with households running down their inventory. We discuss this further in Section 5.1.

Assuming that the decline in quantity purchased reflects households running down inventory,

[9]To abstract from changes in purchase prices around the move, our dependent variable is log quantity purchased. So that we can aggregate quantities across items we restrict attention to goods measured in ounces. This is the majority of goods and includes a wide variety of items. For example, it is over 80% of non-carbonated soft drink UPCs, over 97% of milk UPCs, over 97% of candy UPCs, over 99% of cheese UPCs.

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. Hence, the observed decline is likely much smaller than total inventory.

For households moving to a new 3-digit ZIP Code, the cumulative quantity decline is 9.2% of annual quantity purchased.[10] Restricting the sample to households moving more than 975 kilometers (the top quartile of move distance) the cumulative decline is 15.7% of annual quantity purchased.[11] This provides additional evidence, independent of our assumptions in Section 3, that households hold at least several hundred dollars of inventory.

3.4 Heterogeneity in Inventory Holdings Across Households

In the previous sections, we demonstrate that consumer goods inventory varies substantially across households. Some of this variation may be driven by factors such as taste or unobserved preferences regarding consumption behavior, but others may be correlated with observable characteristics of household demographics, labor force status, and geographic location.

In Table II, we examine the extent to which such observable characteristics of Nielsen households are correlated with their average inventory holdings relative to total consumption. Many of these household characteristics do correlate significantly with inventory holdings.

Older households hold less inventory relative to consumption, as do households who contain a married couple. In contrast, having young children or living in a single family home increases inventory holdings, on average. We also see variation across ethnicity, with white households holding less in inventory and Asian households holding more. Household size actually decreases inventory holdings relative to consumption, but this is primarily driven by an increase in annual consumption, the denominator in the dependent variable.

We note that households in higher priced and denser areas actually hold more inventory relative to consumption. This goes in the opposite direction as predicted, as such households presumably pay higher storage costs than those in low-cost areas. In Appendix Table C.1, we use data from a FinTech company to show that this counterintuitive effect may be also influenced by the 'extensive' margin of food shopping. Households in more expensive locations tend to spend proportionally more on restaurants rather than grocery goods which may prompt differences in the types of goods purchased and the amount of inventory accumulated.

While most of these characteristics are statistically significantly correlated with inventory ratios, the overall explanatory power of even the combination of these variables is very small. The fact that household characteristics explain a low portion of overall shopping behavior is consistent with findings in Hendel and Nevo (2006), who document that using household characteristics to predict the likelihood to purchase a product during a sale yields a similarly low R^2 (under 0.03).

In columns 2-4, we add in two measures that track indicators of product mixes purchased by a given Nielsen household. Including data on the average durability/storability of their product

[10] The U.S. has about 900 3-digit ZIP Codes compared to more than 3,000 counties. By focusing on 3-digit ZIP codes, we therefore reduce the number of local relocations that would permit households to easily move their household inventory. Given the limited number of movers, we extend the sample back to 2006 for this analysis.

[11] Using the imputation estimator proposed by Borusyak, Jaravel and Spiess (2021) gives very similar results (Figure C.3).

mix alongside the average potential bulk savings increases the R^2 somewhat. In particular, average product durability/storability has the largest effect on explanatory power. While choices over products affect inventory accumulation due to inherent characteristics of those products (eg. being easier to store for long periods), identifiable household characteristics do not seem to play a large role in governing inventory accumulation within our sample.

4 A Model of Optimal Household Inventory Management

4.1 Model Overview

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 taking 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 borrowing behavior and portfolio allocation, we need to know the marginal financial return to allocating additional funds to household working capital. In this section, we use the NCP data to calibrate a model of optimal household inventory management. We then use the model to compute the (net) marginal return to household working capital investment, taking into account product depreciation costs and shopping trip fixed costs, 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. Much of that previous literature has exploited temporary price shocks to identify the stockpiling channel of consumer responses to these shocks.

In contrast, we are interested in how allocating a marginal dollar to household working capital facilitates savings. In our model, the household minimizes the cost of providing an exogenously given consumption flow subject to a working capital constraint. For simplicity, we assume that the flow of consumption is constant both between trips and across trips.[12] We define the return to working capital investment as the reduction in cost generated by relaxing the working capital constraint. There are two types of costs: a fixed trip cost and a variable cost per unit purchased. The variable cost reflects both the price charged by the store and the cost of storing the product between purchase and consumption. We assume that this storage cost corresponds to physical quality deterioration and calibrate it using shelf-life data.[13]

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

[12] This assumption can be relaxed. For the CES case, see Baker et al. (2021).

[13] We assume the opportunity cost of funds is zero when solving the model. This means that our computed return to working capital is an absolute rather than a relative return.

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

Second, we model the household's choice of the time interval between trips, Δ (where Δ is measured in months) subject to a working capital constraint. We incorporate the price and working capital functions derived in the first stage. The second stage is similar to standard inventory models (for example, Arrow, Harris and Marschak (1951)), where a firm must decide when to place orders in order to minimize the cost of meeting demand.[14] We refer to the first stage as the 'stockpiling problem' and the second stage as the 'trip-timing problem'. To our knowledge we are the first to integrate these two problems and incorporate a working capital constraint.

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

Our model incorporates two types of savings: buying in larger quantities ("bulk") and buying items on sale ("deals"). In turn, these drive two key relationships between unit prices and shopping trip frequency. Buying in bulk relates directly to the size of the trip (i.e. the amount spent per trip) and buying items on sale relates directly to the frequency of the trip (i.e. more frequent trips yield on average more items on sale for a given trip size). 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. Depending on various parameters (amount of household working capital, depreciation rate, shopping trip fixed cost, frequency and magnitude of sales, etc.), households may prefer one shopping policy over the other.

4.2 Stockpiling problem

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

Although the stockpiling component of our model is similar in some respects to Boizot et al. (2001) and Hendel and Nevo (2006), there are also some crucial differences. We assume that the household observes prices only upon entering the store and there is no fixed cost of making purchases once

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

there. Rather than focusing on a single product, we also assume the household makes purchases to supply an exogenous consumption stream of a continuum of products. Sales rotate across products, but the share of products on sale is the same each trip. The decision to go to the store is not precipitated by a low price realization for a single product of interest. Instead, there is a constant trip interval Δ which we solve for in Section 4.3. Our problem is effectively deterministic after aggregating across products.

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

We require the household to purchase an integer number of 'packs' of each product. One 'pack' is the quantity that must be purchased on the current trip in order to supply consumption until the next trip. We assume constant consumption of all products and so cannot have the household running out of a product between trips.

The household's strategy takes an (s, S) form, where the boundaries depend on the current price p. That is, if the household observes price p and current inventory is at or below $s(p)$, the household makes a purchase and raises inventory to level $S(p)$. The optimality of this strategy is proven for the continuous case by Hall and Rust (2007) and Sethi and Cheng (1997) and it is also shown there that $s(p)$ and $S(p)$ are decreasing in p.

Intuitively, when the price is at its maximum level p^f, it does not make sense to buy more than is required for consumption before the next shopping trip because the price next trip cannot be any higher (and may be lower). Therefore, the household only makes a purchase at full price when it has no product left in stock. That is, $s(p^f) = 0$ and $S(p^f) = 1$.

In contrast, when a product is on sale, the current price is lower than the *expected* price at the time of the next trip (i.e. $p^d < xp^d + (1-x)p^f$). Depending on depreciation costs, it may make sense for the household to buy more than it requires for consumption in the current trip interval. Because p^d is the lowest possible price and there is no fixed cost of purchasing, $S(p^d) = s(p^d) + 1 \geq 0$ (note that one pack is consumed over the trip interval). Therefore, the only aspect of the policy we still need to solve for is $s(p^d)$. For simplicity, we denote $s(p^d)$ by m from now on and refer to m as the stockpiling strategy. The order quantity q^o the household purchases when observing p^d is given by $q^o = m + 1 - q$, where q is the current quantity in stock. The household purchases one pack for immediate consumption and $m - q$ packs for storage. Therefore, m is also the maximum number of packs of that product the household will purchase for storage (or equivalently the number of trips they are willing to buy in advance of consumption).

4.2.1 Expected price conditional on the stockpiling strategy

We would like to find an expression for the long-run quantity-weighted expected price paid conditional on the stockpiling strategy m, $E[p|m]$. To do this, we need to know the share of sale purchases for a given choice of m. Let $x_{ss}(m)$ be the steady-state share of purchases that occur at p^d. The long-run expected price paid conditional on strategy m is then:

$$E[p|m] = \frac{p^d x_{ss}(m) E[q^o|\text{purchase at } p^d, m] + p^f (1 - x_{ss}(m)) E[q^o|\text{purchase at } p^f]}{x_{ss}(m) E[q^o|\text{purchase at } p^d, m] + (1 - x_{ss}(m)) E[q^o|\text{purchase at } p^f]} \tag{7}$$

We can simplify this as we know that conditional on a purchase occurring at p^f, exactly one pack will be purchased. This gives us:

$$E[p|m] = \frac{p^d x_{ss}(m) E[q_o|\text{purchase at } p^d, m] + p^f (1 - x_{ss}(m))}{x_{ss}(m) E[q^o|\text{purchase at } p^d, m] + (1 - x_{ss}(m))} \tag{8}$$

Next we will find expressions for $x_{ss}(m)$ and $E[q^o|\text{purchase at } p_d, m]$.

Under the assumptions stated above, the price at which a transaction occurs follows a first-order Markov process. Note by restricting attention to the set of times at which transactions occur, we avoid having to include the level of inventory as a state variable. This is because conditional on a transaction occurring, the past level of inventory becomes irrelevant (at a transaction time we know the current level of inventory must be $m + 1$ packs if $p = p^d$ and one pack if $p = p^f$).

The transition matrix is below. The first row (column) of the matrix corresponds to p^f and the second row (column) to p^d. The rows correspond to the price at which transaction t occurs and the columns correspond to the price at which transaction $t + 1$ occurs:

$$X_{ss} = \begin{pmatrix} 1 - x & x \\ (1 - x)^{m+1} & 1 - (1 - x)^{m+1} \end{pmatrix}$$

The intuition is as follows. When a transaction has just occurred at p^f there is currently one pack in stock and a transaction must occur on the following trip. As prices are iid, the probability the next transaction occurs at p^d is x and the probability it occurs at p_f is $1 - x$.

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

Now we have the transition matrix, solving for the steady state distribution is straightforward. As we restrict attention to times at which transactions occur, the solution will be the steady state distribution of transaction prices. This is what we need in order to compute the expected price paid. The steady state deal probability x_{ss} solves:

$$\begin{pmatrix} 1 - x_{ss} & x_{ss} \end{pmatrix} \left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} 1 - x & x \\ (1 - x)^{m+1} & 1 - (1 - x)^{m+1} \end{pmatrix} \right] = \begin{pmatrix} 0 & 0 \end{pmatrix}$$

Simplifying:

$$\begin{pmatrix} 1-x_{ss} & x_{ss} \end{pmatrix} \left[\begin{pmatrix} x & -x \\ -(1-x)^{m+1} & (1-x)^{m+1} \end{pmatrix} \right] = \begin{pmatrix} 0 & 0 \end{pmatrix}$$

This gives:

$$x_{ss} = 1 - \frac{(1-x)^{m+1}}{x+(1-x)^{m+1}} = \frac{x}{x+(1-x)^{m+1}} \tag{9}$$

Next, we need to work out the steady-state average quantity purchased conditional on a sale, $E[q^o|\text{purchase at } p^d, m]$. Following a transaction at p^d the household has $m+1$ units in stock. For each additional period when a sale is not observed, the household adds one extra pack upon next observing p^d. That is, if p^d is observed again on the next trip the household buys one unit. If p^d is next observed after two trips the household buys two units, and so on. The probability that t periods pass without a sale being observed is $(1-x)^t$. The expected quantity purchased conditional on a transaction occurring at p^d is therefore:

$$E[q^o|\text{purchase at } p^d, m] = \sum_{t=0}^{m} (1-x)^t \tag{10}$$

Substituting Equations 9 and 10 into Equation 8 gives:

$$E[p|m] = \frac{p^d \frac{x}{x+(1-x)^{m+1}} \sum_{t=0}^{m}(1-x)^t + p^f \frac{(1-x)^{m+1}}{x+(1-x)^{m+1}}}{\frac{x}{x+(1-x)^{m+1}} \sum_{t=0}^{m}(1-x)^t + \frac{(1-x)^{m+1}}{x+(1-x)^{m+1}}} \tag{11}$$

Simplifying:

$$E[p|m] = \frac{p^d x \sum_{t=0}^{m}(1-x)^t + p^f (1-x)^{m+1}}{x \sum_{t=0}^{m}(1-x)^t + (1-x)^{m+1}} \tag{12}$$

Because $x\sum_{t=0}^{m}(1-x)^t + (1-x)^{m+1} = 1$:

$$E[p|m] = p^d \left(1 - (1-x)^{m+1}\right) + p^f (1-x)^{m+1} \tag{13}$$

Equation 13 can also be obtained more directly using an intuitive argument. The only case where the household pays p^f in the long-run is when there have been $m+1$ sequential no sale trips (and only one pack is bought in that case). In all other cases the household pays p^d. The formal derivation is helpful, however, for two reasons. First, we have so far ignored depreciation costs. To accurately incorporate these costs it is necessary to keep track of the distribution of when deal purchases were made. The distribution of purchase times is also necessary for computing the level of inventory implied by strategy m.

4.2.2 Incorporating depreciation into the expected price

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

There are two equivalent ways to think about the depreciation cost. We can convert the price paid in store to an 'effective' price, which reflects quality deterioration. Alternatively, we can think of

the depreciation cost as modifying the quantity that must be purchased to meet consumption needs: if items randomly go bad, more needs to be purchased initially. As the total price paid for good i on a trip is $p_i q_i^o$, it does not matter whether we apply the depreciation cost factor to the price or the quantity.

If a product deteriorates exponentially at rate δ, a unit purchased at time zero for consumption following trip t effectively costs $e^{\delta \Delta t}$ times more than a product purchased for immediate consumption. To adjust the expected price function (Equation 13) to account for these costs, we need to consider the whole distribution of purchase times for a product purchased on sale. It is convenient to incorporate depreciation costs into our expression for $E[q^o|\text{purchase at } p^d, m]$. This gives:

$$E[q^o|\text{purchase at } p^d] = \sum_{t=0}^{m} e^{\delta \Delta (m-t)} (1-x)^t$$

The expression for the conditional expected price is then:[15]

$$E[p|m] = p^d x \sum_{t=0}^{m} e^{\delta \Delta (m-t)} (1-x)^t + p^f (1-x)^{m+1} \tag{14}$$

As discussed above, the intuition for this expression is that for each additional trip when a sale is not observed, the household adds an extra pack upon next observing p_d. As we increase the number of trips without a sale, t, the depreciation cost of the incremental pack ($e^{\delta \Delta (m-t)}$) declines because that pack will not be held for as long prior to consumption.

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

$$\Delta \cdot (m+1) \leq \bar{t} \tag{15}$$

This is intended to make the model more realistic. While many products do deteriorate gradually from the date of purchase, there is also typically a point at which quality falls below the minimum required for consumption. The expiration date of the product provides guidance to the consumer regarding when this point has been reached, and this is ultimately how we will calibrate $\bar{t}$. The method for determining expiration dates depends on the product. For some products, it is a question of food safety. For others, expiration dates are driven by perceived quality falling below a cutoff. Regardless of the reason, we think it is unlikely that households would set m so high that they would be unable to consume the item prior to expiration.

4.2.3 Aggregate inventory

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

[15] Depreciation costs incurred *during* the trip interval will be accounted for in the trip-timing problem. For this section, it helps to think of consumption as occurring exactly at the time of each trip.

constraint to compute the marginal net return to household working capital.

Intuitively, the higher the value of m, the higher the level of inventory will be when going to the store. If the household does not stockpile at all (i.e. q^o is always equal to 1), inventory is exactly zero at the time of the next trip.

We first derive an expression for the number of packs in stock immediately prior to a trip, $I(m)$. As these packs were all purchased on sale, we value them at price p^d. To get the value of inventory following each trip, we add the value of purchases made on each trip. This is discussed in Section 4.3. Here, we focus on $I(m)$.

We aggregate over the continuum of goods i consumed by the household, assuming that sales are iid both across goods and over time. The amount of good i left in stock prior to the trip depends on when a sale was last observed and the household's choice of m. In general, the household has $m-t$ packs in stock of share $x(1-x)^t$ of products immediately prior to each trip. For example, $t=0$ corresponds to goods for which a sale was just observed last trip. On each trip the household observes a sale for share x of goods and following the trip has m packs of these goods in stock. Integrating across all products i, the total number of packs in stock immediately prior to any trip is:

$$I(m) = \mathbb{1}_{\{m>0\}} x \sum_{t=0}^{m-1} (m-t)(1-x)^t \tag{16}$$

4.3 Trip-timing problem

Next, we integrate the price equation (14), shelf-life constraint (15) and the inventory equation (16) into the trip-timing problem. The household will choose both m and Δ to minimize the cost of providing an exogenous monthly consumption stream, subject to a working capital constraint. The cost per trip can be decomposed into two components – a fixed cost (e.g. the opportunity cost of time spent shopping or pecuniary costs of travel) and a variable cost which depends on the quantity of products purchased.

We divide products into perishability groups indexed by l. Monthly consumption of perishability group l is C_l. We use k to denote the fixed trip cost and P_l to denote the per unit price of group l. P_l will depend on the trip interval Δ and also on the household's stockpiling strategy, $\{m_l\}$.[16] In addition to stockpiling in response to temporary deals, households can also save by buying larger pack sizes and paying a lower per unit price. Larger trip sizes correspond empirically to a household either consuming more or shopping less frequently. We incorporate bulk discounts by multiplying the expected price function from Section 4.2.2 by a bulk discount function $b(Q_l)$:

$$P_l(\Delta, m_l) = b(Q_l) E[p|m_l, \Delta]$$

where $b(Q_l) > 0$, $b'(Q_l) < 0$ and $b''(Q_l) > 0$. We discuss the calibration of $b(Q_l)$ in Section 4.5.4. Note that the expected price is achieved with certainty and $P_l(\Delta, m_l)$ is deterministic.

The household must purchase enough of each good to last until the next trip, and we previously defined a 'pack' as containing exactly this amount.[17] Next, we work out the quantity per pack. For

[16] We allow the household to make a separate stockpiling choice, m_l, for each perishability group l. So far we abstracted from this for simplicity. Each perishability group l contains a continuum of individual products i.

[17] We do not allow households to set different values of Δ for different goods. Although setting different values of Δ

a good that deteriorates exponentially at a rate δ_l, the quantity that must be purchased to satisfy continuous monthly consumption flow C_l over the trip interval Δ is:

$$Q_l(\Delta) = \int_0^{\Delta} e^{\delta_l t} C_l dt = \frac{C_l}{\delta_l}(e^{\delta_l \Delta} - 1) \tag{17}$$

That is, when the time between trips is Δ and the household buys quantity $Q_l(\Delta)$ each trip, inventory next hits zero precisely when the next shopping trip is scheduled to occur. Intuitively, if $\delta_l = 0$ the amount that needs to be purchased is simply ΔC_l (i.e. the trip interval in months multiplied by the monthly consumption flow). To the extent that product quality declines as time since purchase increases, more needs to be purchased initially.

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

$$\frac{k + \sum_l P_l(\Delta, m_l) Q_l(\Delta)}{\Delta}$$

Next, we need to incorporate the working capital constraint. In Section 4.2.3, we derived $I(m)$, the number of packs in stock immediately prior to a trip (Equation 16). The amount of working capital required to facilitate a choice of $\{m_l\}$ and Δ is the maximum inventory over the trip cycle, which occur immediately following a trip. At this point in time 100% of household working capital is held as stored inventory goods. Between trips, the inventory share of working capital gradually declines through consumption and depreciation and is replaced with accumulated cash used to pay for the next trip.

As all stockpiling occurs at price p^d, the inventory of good l prior to a trip is $I_l(m_l)Q_l(\Delta)p_d$. This is the number of packs in stock, times the pack size in number of consumption units, times the price per consumption unit. To compute the working capital required to support $(\Delta, \{m_l\})$, we need to add the value of products purchased on a single trip, $P_l(\Delta, m_l) \cdot Q_l(\Delta)$. The maximum allowable level of working capital is exogenous and is denoted by $\bar{I}$. We will compute the marginal return to working capital by gradually increasing $\bar{I}$.

We now have all the elements we need. The household minimizes the discounted cost (18) of providing the exogenous consumption flow C_l subject to a working capital constraint (19) and a restriction on storage time for each perishability level l, i.e. shelf life constraints (20):

$$V(\bar{I}) = \min_{\Delta, \{m_l\}} \frac{k + \sum_l P_l(\Delta, m_l) Q_l(\Delta)}{\Delta} \tag{18}$$

$$s.t. \quad \sum_l P_l(\Delta, m_l) Q_l(\Delta) + I_l(m_l) Q_l(\Delta) p_d \leq \bar{I} \tag{19}$$

$$\Delta \cdot (m_l + 1) \leq \bar{t}_l \; \forall \; l \tag{20}$$

Ultimately, we are interested in the relationship between the dollar amount invested in household working capital, $\bar{I}$, and cost of providing the consumption stream. In order for a particular

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).

shopping strategy to be feasible, the level of inventory immediately following a trip must not exceed the amount of household working capital $\bar{I}$. We will solve the problem for different levels of $\bar{I}$, and use this to compute the return to "investing" in household working capital (i.e. marginally increasing $\bar{I}$). The investment payoff will be the reduction in the cost V, so that we can define the marginal (net) return to household inventory management as $r_I(\bar{I}) = V'(\bar{I})$.

4.4 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 and shelf life constraints are satisfied and find the combination that minimizes the cost function.

4.5 Calibration

We calibrate the model by choosing parameters to match a number of data moments. Consumption flows C_l and trip length Δ are expressed in months. For example, $\Delta = 0.23$ implies a trip interval of 0.23 months (i.e. 1 week). Similarly $\sum_l C_l = 1$ implies that the household consumes a total of one unit over the course of one month. Depreciation costs $\{\delta_l\}$ are monthly rates with continuous compounding. For example, if $\delta = 0.5$, each unit purchased corresponds to $e^{-0.5} = 0.61$ units one month later.

To assist with calibrating the bulk discount function $b(\Delta)$, we define a standard trip size ($\hat{Q}_l$) as the trip size corresponding to the optimal choice of Δ in the absence of bulk discounts.

The problem scales with consumption if we assume that the trip cost k and standard trip size $\hat{Q}_l$ both scale with consumption. We therefore set $\sum_l C_l = 1$ and express the trip cost and standard trip size as shares of monthly consumption. We set C_l equal to the expenditure share of group l.

4.5.1 Fixed trip cost

We set the fixed cost per shopping trip so that the trip interval in the model matches the average trip interval between trips to grocery or discount stores in the Nielsen data. The average trip interval across households is 0.25 months (or roughly one week) and the corresponding fixed cost per trip is 0.57% of monthly consumption.

4.5.2 Depreciation

We start by manually assigning a shelf life for each Nielsen product group using information on product life from FSIS. We then allocate each Nielsen product group to one of four perishability groups. For each perishability group l, we will calibrate a depreciation rate δ_l and a shelf-life $\bar{t}_l$.

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

In contrast, variation in shelf-life is important for more perishable items. We therefore try to choose the cutoffs for the remaining groups sensibly so that we can capture the variation well with a small number of groups. We use a decision tree to divide the set of products with shelf-life less than 6 months into 3 subgroups. Our approach selects cutoffs so that the products within each subgroup are as similar as possible with respect to their log shelf-life.[18] The most perishable group ($l = 1$) comprises fresh meat, fresh baked goods, and pre-prepared salads. The remaining perishability groups contain a larger number of Nielsen product groups. Milk and yoghurt are examples of items in Group 2. Eggs, packaged deli meats and fruit juice are in Group 3. Group 4 contains all product groups with a shelf-life of more than 6 months, and includes items such as cereal, carbonated beverages, cleaning products, toiletries, canned items and dried grains and vegetables.

When mapping shelf life to the model depreciation relationship we need to account for differences in the depreciation profile across product types. In Appendix Table C.2 we provide a description of how quality changes over time for a number of different food items. For items such as coleslaw and bread, quality measures deteriorate rapidly from the date of purchase. In line with this, we assume exponential depreciation for more perishable products. For Groups 1-3 we calibrate δ_l so that the consumption value on the expiration date is 50% of the value on the purchase date and then set $\bar{t}_l$ equal to the time until expiration.[19]

In contrast, Table C.2 shows that for more storable products, such as breakfast cereal, snack bars and shelf-stable ready meals there is effectively no decline in quality for the first few months following purchase. In addition, the main limiting factors for these products are things like flavor and texture changes, rather than the item becoming unsafe (Singh, 1994). We therefore believe that exponential depreciation is inappropriate for these products. Instead, we assume that the product does not depreciate prior to the expiration date, at which point the product is disposed of (that is we set $\delta_l = 0$ and require $(m_l + 1) \cdot \Delta \leq \bar{t}_l$).[20]

Table III shows the calibrated values of δ_l and $\bar{t}_l$ for each each perishability group along with expenditure shares s_l. The least perishable group is by far the largest, accounting for 71% of NCP expenditure. We want to ensure that $\bar{t}_l$ is at least as long as the data trip interval for all groups. We therefore set $\bar{t}_1 = \Delta$, which implies that $m_1^* = 0$. This is a fairly minor adjustment as the average shelf life in group 1 is 0.16 months, compared with a data trip interval of 0.25. In practice, the exact value of $\bar{t}_1$ is in any case uncertain. Some products may be stored for a bit longer than the standard shelf

[18] We use log shelf-life to group based on the log difference rather than the absolute difference in shelf-life. A given absolute difference in shelf-life is less important for products with a longer shelf-life for the reasons discussed above.

[19] For each perishable Nielsen product group g with expiration date in months E_g and expenditure share (of perishable products) s_g, we compute the depreciation rate δ_g which leaves 50% of the product remaining on the expiration date (i.e. we assume the half life is equal to the time to expiration). The perishable consumption value remaining after one month is $\sum_g s_g e^{-\delta_g}$, and δ_1 solves $e^{-\delta_1} = \sum_g s_g e^{-\delta_g} = \sum_g s_g e^{\frac{1}{E_g}\log 0.5} = \sum_g s_g 0.5^{\frac{1}{E_g}}$ such that $\delta_1 = -\log\left(\sum_g s_g 0.5^{\frac{1}{E_g}}\right)$. Depreciation could be a deterministic change in quality, but may also have a stochastic interpretation. For example, the life of an item after purchase may vary depending on storage time and conditions prior to purchase. Mapping the expiration date to a depreciation rate is subjective. In practice different cutoffs and metrics may be used for different products. Alternative approaches, such as assuming the expiration date corresponds to the mean expiry time, give qualitatively similar results.

[20] Allowing for gradual depreciation to begin a few months after purchase but prior to expiration would arguably be more realistic, but would have little effect on our results. As discussed above, any variation in product depreciation more than a few months out has little effect on the return to working capital because the price reduction from stockpiling is almost fully exploited by this point.

life, or consumed close to the trip time rather than continuously.[21]

4.5.3 Price Process

To calibrate p^f and p^d, we estimate the average price drop associated with a discount event in the NCP for each perishability group. We 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 following regression for each perishability group l:

$$\log \text{Price}_{u,t} = \beta_1 D_{u,t} + \beta_2 \text{Bulk}_{u,t} + \text{Month FE}_t + \text{Product}_u \times \text{Household FE}_i + \epsilon_{u,t}$$

where u indexes transactions at the UPC level. We then convert the log price difference into a price ratio. For example, for Group 1 we estimate a log price difference of 0.3 and we then set $\frac{p_1^f}{p_1^d} = e^{0.3} = 1.35$. We pin down p_l^f and p_l^d by normalizing the expected price achieved using a non-strategic shopping policy without bulk discounts, i.e. $E[p_l] = xp_l^d + (1-x)p_l^f = 1$. p_l^d and p_l^f are similar across perishability groups.

To calibrate the probability of a sale, x, we estimate the following relationship using the NCP separately for each perishability group l:

$$\text{Deal \& Coupon Share}_{h,g,y} = \alpha + \beta_1 \text{Inventory Ratio}_{h,g,y} + \beta_2 \text{Trips}_{h,g,y} + \epsilon_{h,g,y} \tag{21}$$

where $\text{Deal \& Coupon Share}_{h,g,y}$ is the share of items purchased by household h in product group g in calendar year y where either a coupon was used or the item was reported to be on sale. $\text{Inventory Ratio}_{h,g,y}$ is household h's average inventory of product group g in year y (as described in Section 3) divided by household h's total Nielsen spending on product group g in year y.

We set x equal to the predicted deal and coupon share for a household with an inventory ratio equal to $0.5/\text{Trips}_{h,g,y}$, where $\text{Trips}_{h,g,y}$ is the number of trips on which household h purchased product group g in year y. This is the inventory ratio a household would have in the model if they did not engage in strategic deal shopping. The calibrated value of x ranges between 0.2 and 0.26, with more storable products experiencing deals slightly more often. This means the discounted price is observed about every 4$^{\text{th}}$ or 5$^{\text{th}}$ trip. The calibrated price process parameters, p_l^d, p_l^f and x^l are shown in Table IV.

4.5.4 Bulk discount function

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

$$b(Q_l) = \alpha + \beta e^{-\sigma \frac{Q_l}{\bar{Q}_l}} \tag{22}$$

[21] This type of behavior is a realistic violation of the constant consumption assumption we make in Section 3. However, as we discuss in Appendix B, violations similar to this one would not have a substantial effect on the inventory calculation (this particular example is somewhere between Figure C.7a and Figure C.7b).

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

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 effective price per unit in the absence of discounts to $P_l(\hat{\Delta}, 0) = 1$ (i.e. with $m_l = 0$ and $b(Q_l) = 1 \forall \; Q_l$). The parameter α is interpreted as one minus the maximum % savings that can be obtained from buying in bulk.

To calibrate the bulk discount function, we first need to 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.[22]

To standardize prices and pack size, we compute the average number of units in the second quintile of the pack size distribution for each product, as well as the expenditure weighted annual average price per unit. We then divide by these second quintile averages.[23] Motivated by the observation that the first quintile appears to contain travel size packs which are substantially more expensive on a per unit basis, we assume the second quintile of pack size corresponds to the "standard" trip size in the model.[24]

In our model, buying larger pack sizes corresponds to a reduction in trip frequency (holding consumption fixed). Therefore, quantities of all products are scaled up in the same way. In the data, a large proportion of spending is accounted for by products that 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). Estimating the price-quantity relationship using the raw data will therefore overstate the bulk savings households could achieve with respect to their total consumption. Instead, we group standardized pack sizes into bins and expand the dataset so it is balanced. We then fill in prices for missing pack size bins using prices from neighboring bins.[25]

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

$$\text{Price}_{p,q} = a_0 + a_1 e^{-\sigma \text{Units}_{p,q}} \tag{23}$$

[22] We 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 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.

[23] We exclude items measured in "count" (CT) as it is unclear whether the units are comparable across UPCs.

[24] In some cases no UPCs are allocated to the second quintile due to insufficient variation in pack size (and given the constraint that UPCs with the same pack size cannot be allocated to different quintiles). For these items we treat the first quintile as the standard pack size.

[25] When a pack size bin does not contain any UPCs associated a particular product, we assign the price from the nearest bin containing any UPCs for that product. If there are bins equally close either side we take the average of the two prices.

where $\text{Price}_{p,q}$ is the standardized unit price of product p at pack size q and $\text{Units}_{p,q}$ is the standardized number of units of product p at pack size q. We then choose σ to maximize the within-R^2. We perform this procedure separately for each level of l. Columns 1-3 of Table V show the the values of σ and estimates of a_0 and a_1 for each group. In the model, we normalize the price of the standard pack size to one, and the price of other pack sizes reflect percentage deviations from the standard pack size. We therefore calibrate α and β using $\alpha = \frac{\hat{a}_0}{\hat{a}_0+\hat{a}_1 e^{-\hat{\sigma}}}$ and $\beta = \frac{\hat{a}_1}{\hat{a}_0+\hat{a}_1 e^{-\hat{\gamma}}}$. The values of α and β for each perishability group are shown in Columns 4 and 5. Figure IV compares the bulk discount functions we use in the model with the corresponding relationships in the data.

4.6 Implications for Portfolio Choice

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

We consider the effect of working capital on the cost of supplying consumption to be analogous to interest earned on an investment. For simplicity of exposition, we assume that both trip fixed cost and bulk savings scale proportionally to annual consumption so that the portfolio choice problem can be solved independently of the consumption problem. This assumption is reasonable in cases where the opportunity cost of time is increasing in consumption (or in "permanent income"), and where increases in consumption are reflected in purchases of higher quality products rather than purchasing larger quantities.

Consider a household with 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-period wealth (or consumption) by solving the following problem:

$$\max_{\alpha_I,\alpha_f,\alpha_m} EU\big((1+\tilde{r}_p)w\big) \tag{24}$$

$$s.t. \quad \tilde{r}_p = \frac{1}{w}\int_0^{\alpha_I w} r_I(x)dx + \alpha_f r_f + \alpha_m \tilde{r}_m \tag{25}$$

$$1 = \alpha_I + \alpha_f + \alpha_m \tag{26}$$

where α_I is the share of initial wealth w allocated to working capital (so $\bar{I} = \alpha_I w$) with marginal return r_I, α_f is the share allocated to the risk-free bond with return r_f, and α_m is the share allocated to the risky asset with stochastic return $\tilde{r}_m$.[26] We utilize a market return that is common across households $\tilde{r}_m$ (e.g. the mean-variance efficient return), but note that other work has shown that returns obtained by lower income households typically under-perform an optimal portfolio and would thus

[26]The average inventory portfolio share in Figure II is closely related to α_I. α_I additionally includes a cash component, but this is a small share of (model) working capital over the range of inventory ratios we observe in the data.

further drive down demand for the risky asset (e.g. Calvet, Campbell and Sodini (2007)).

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

Assuming consumers are risk averse, they choose $\alpha_I = 1$ as long as the marginal return to working capital investment $r_I(w) \geq E[\tilde{r}_m]$ 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 $r_I(w) < E[\tilde{r}_m]$ the optimal allocation depends on the utility function, but as long as $r_I(w) > r_f$ 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. Consequently, α_I gradually declines as wealth increases.

5 Financial Net Returns to Household Inventory Investment

Solving the optimization problem (18) 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 both the price paid in store and also 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. It therefore makes sense to consider a somewhat larger increment. When reporting our results we use an increment of 2.5% of annual consumption. We multiply by 12 to convert monthly to annual returns.[27]

Table VI shows how increasing the maximum household working capital $\bar{I}$ affects 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_l for storable products.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 smaller and more 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. By shopping more frequently, households can buy a larger share of

[27] We assume that working capital 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 we assume.

items on sale at a given level of working capital.

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 pushes the trip interval up and raises bulk savings. This effect is strongest at very low levels of working capital. Given that we match the average NCP trip interval of about one week, a relatively small amount of working capital is required to achieve the desired trip interval.

In general, model outcomes such as the marginal return, percentage of working capital held as 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 stockpile items on sale, this can put downward pressure on trip size due to depreciation costs and reduce bulk savings.

At low levels of household working capital investment, the marginal return to additional investment is very high. When household working capital is equal to 5% of annual consumption, the marginal return is around 60%. The marginal return gradually diminishes and reaches zero when household working capital is around one third of annual consumption. As we discuss in Section 4.6, households in the model do not participate in the stock market if the marginal return to working capital is more than the expected stock market return. Over the two decades prior to our sample, the average annual S&P 500 return was around 8%. In Table VI, this corresponds to working capital and inventory share cutoffs of around 12.5%-15% of annual consumption. When aggregating to the product group level, around 15% of households have an inventory ratio below this cutoff. The share increases to 75% when aggregating to department. In practice we expect this cutoff could vary substantially across households depending on their financial education and investment opportunities. In addition, some share of inventory holdings is explained by factors not considered in our model. For example, goods stockpiled for emergencies. To the extent that these other factors raise the value of inventories, the non-participation cutoff for the inventory ratio could be higher than 15%.

As working capital increases, the share allocated to inventory increases and the share allocated to cash declines. This is because additional working capital is increasingly used to stockpile goods on sale rather than to increase trip size. 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 VII we compute returns for households who have a very large trip fixed cost of 10% of monthly consumption. 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 at

low levels of 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 VII 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. Allowing for substitution away from these perishable products when trip fixed costs rise would lead to 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.[28] Secondly, we use the average number of days between a household's trips as a proxy for their unobserved 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 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 and NRP.[29]

5.1.1 Deal Savings

We measure deal savings as the discount relative to the amount the household would have spent if they paid the average UPC-ZIP3 price over that year. That is, deal savings represent the difference between the average price of a UPC and the price the household actually pays: the *additional* savings that resulting from strategic shopping behavior 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. We discuss the construction of the deal savings measure in more detail in Appendix A.

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. This is consistent with

[28] 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.

[29] Because our deal and bulk savings measures require us to also use the NRP as well as the NCP, we restrict attention to UPCs that appear in both datasets. We also recompute our inventory measure using the approach described in Section 3. The ratio of inventory to spending is somewhat higher in this sample than it is the full sample, indicating that products in both datasets are stockpiled in greater quantities than products in the NCP only.

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 model relationship between savings and the ratio of inventory to annual spending. 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.

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 also observe a positive relationship between inventory and deal savings in the data. 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). 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, preferences for perishable products, or in the way grocery items are processed and stored. In the model, variation in δ_l generates a similar relationship between inventory and deal savings as variation in $\bar{I}$. Variation in holding costs across households may therefore also contribute to the data relationship in Figure Vd.

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 by about 1 percentage point in the months prior to the move and recover gradually in the months following the move.[30]

5.1.2 Bulk Savings

Our measure of bulk savings compares the average unit price of the UPC actually purchased with the average unit price of a second quintile pack size of the same product in the same year and 3-digit Zip Code. Note that by using the average prices the measure abstracts from temporary deal savings. We describe our method for computing bulk savings in more detail in Appendix A.

[30] This result is also robust to using the imputation estimator proposed by Borusyak et al. (2021) (Figure C.4).

Figure VIa plots the model relationship between trip interval and bulk savings. In the model, households with high fixed costs, and therefore longer trip intervals, obtain more bulk savings by assumption. This is because households with less frequent trips also buy more each trip, holding consumption fixed. While the data relationship is also positive, the slope is close to zero (Figure VIb).

There is also a broadly flat relationship between the inventory ratio and bulk savings in the model (Figure VIc). This follows from our assumption that bulk savings are obtained by choosing a larger steady state trip size. Consequently, variation in bulk savings generated by adjusting $\bar{I}$ comes purely through the effect on the trip interval, which we controlled for in this specification for the reasons outlined in Section 5.1.1. Figure VId shows the relationship in the data is similar, supporting our assumption. That is, in both the data and the model, stockpiling mostly reflects households taking advantage of temporary deals on standard pack sizes, rather than buying in bulk. 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. Large packs may instead be appealing to households who consume the product at a high rate.

5.1.3 Total Savings

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:

$$\text{Gross return} = \frac{\text{Dollar in-store savings}(\bar{I}_1) - \text{Dollar in-store savings}(\bar{I}_0)}{\bar{I}_1 - \bar{I}_0} \tag{27}$$

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 annual consumption and plot the relationship. We also plot the relationship in the NCP between total savings 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 a strongly increasing relationship 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 man-

agement 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.

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., Lorenz Kueng, Leslie McGranahan, and Brian T. Melzer, "Do household finances constrain unconventional fiscal policy?," *Tax Policy and the Economy*, 2019.

_ , Stephanie Johnson, and Lorenz Kueng, "Shopping for Lower Sales Tax Rates," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 209–50.

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.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," *Working Paper*, 2021.

Brocklehurst, T.F., "Delicatessen salads and chilled prepared fruit and vegetable products," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, chapter 6, pp. pp87–123.

Broda, Christian and David E. Weinstein, "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 2010, *100* (3), 691–723.

Browning, Martin and Thomas Crossley, "The Life-Cycle Model of Consumption and Saving," *Journal of Economic Perspectives*, 2001, *15*, 3–22.

_ and _, "SHOCKS, STOCKS, AND SOCKS: SMOOTHING CONSUMPTION OVER A TEMPORARY INCOME LOSS," *Journal of the European Economic Association*, 2009, *15*, 1169–1192.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 2007, *115* (5), 707–747.

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 Cyclicality 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.

Corrigan, Virginia, Duncan Hedderley, and Winna Harvey, "Modeling the shelf life of fruit-filled snack bars using survival analysis and sensory profiling techniques," *Journal of Sensory Studies*, 2012, *27*, 403–416.

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.

Duyvesteyn, W. S., E. Shimoni, and T. P. Labuza, "Determination of the End of Shelf-life for Milk using Weibull Hazard Method," *LWT - Food Science and Technology*, May 2001, *34* (3), 143–148.

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).

Gauchez, Hélène, Anne-Laure Loiseau, Passcal Schlich, and Christophe Martin, "Impact of aging on the overall liking and sensory characteristics of sourdough breads and comparison of two methods to determine their sensory shelf life," *Journal of Food Science*, 2020, *85* (10), 3517–3526.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, "Money Doctors," *Journal of Finance*, 2015, *70* (1), 91–114.

Goddard, M.R., "The storage of thermally processed foods in containers other than cans," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, pp. 256–274.

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.

Hall, George and John Rust, "The (S,s) policy is an optimal trading strategy in a class of commodity price speculation problems," *Economic Theory*, 2007, *30*, 515–538.

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** _ , "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.

Howarth, J.A.K., "Ready-to-eat Breakfast Cereals," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, pp. 235–254.

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, 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.

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.

Lewis, M. and R.H. Hale, "Chilled Yogurt and Other Dairy Desserts," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, pp. 127–154.

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).

Orhun, A. Yeşim and Mike Palazzolo, "Frugality Is Hard to Afford," *Journal of Marketing Research*, 2019, *56*, 1–17.

Parker, Jonathan, "The Reaction of Household Consumption to Predictable Changes in Payroll Tax Rates," *American Economic Review*, 1999, *89*, 959–973.

Petersen, Mitchell A. and Raghuram G. Rajan, "Trade Credit: Theories and Evidence," *Review of Financial Studies*, 1997, *10* (3), 661–691.

Rampini, Adriano, "Financing Durable Assets," *American Economic Review*, 2019, *109*, 664–701.

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.

Sethi, Suresh and Feng Cheng, "Optimality of (s, S) Policies in Inventory Models with Markovian Demand," *Operations Research*, 1997, *45* (6), 931–939.

Singh, R.P., "Scientific principles of shelf life evaluation," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, pp. 3–24.

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).

Symons, H., "Frozen foods," in C.M.D. Man and A.A. Jones, eds., *Shelf Life Evaluation of Foods*, 1 ed., Chapman and Hall, 1994, pp. 296–315.

Terpstra, M.J., L.P.A. Steenbekkers, N.C.M. de Maertelaere, and S. Nijhuis, "Food storage and disposal: consumer practices and knowledge," *British Food Journal*, 2005, *107* (7), 526–533.

Vissing-Jørgensen, Annette, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 2002, *110* (4), 825–853.

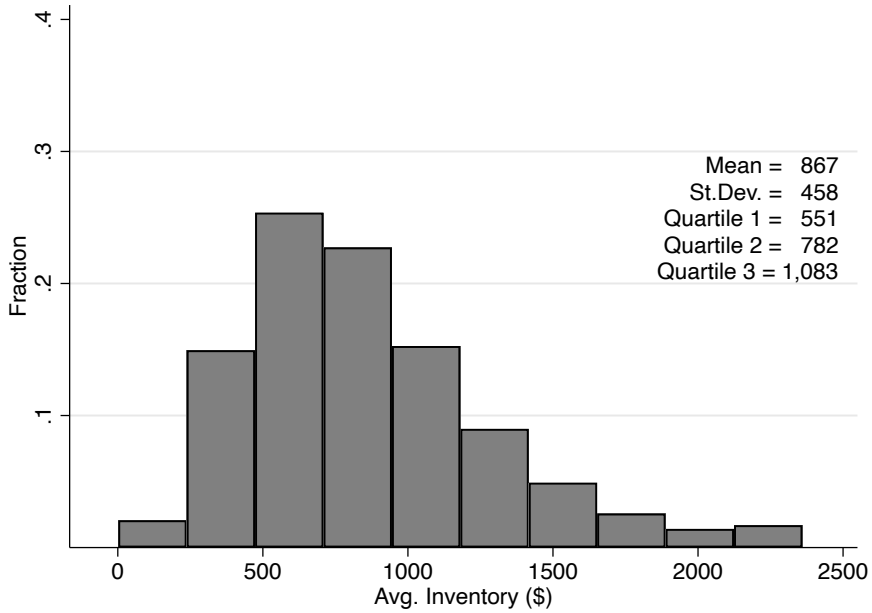
Winger, R.J., "Storage life and eating quality of NZ frozen lamb: A compendium of impressive longevity," in P. Zeuthen, J.C. Cheftel, C. Eriksson, M. Jul, H. Leniger, P. Linko, G. Varela, and G. Vos, eds., *Thermal Processing and Quality of Foods*, Elsevier Applied Science, 1984.

Yang, S. Alex and John R. Birge, "Trade Credit, Risk Sharing, and Inventory Financing Portfolios," *Management Science*, 2018, *64* (8), 3667–3689.

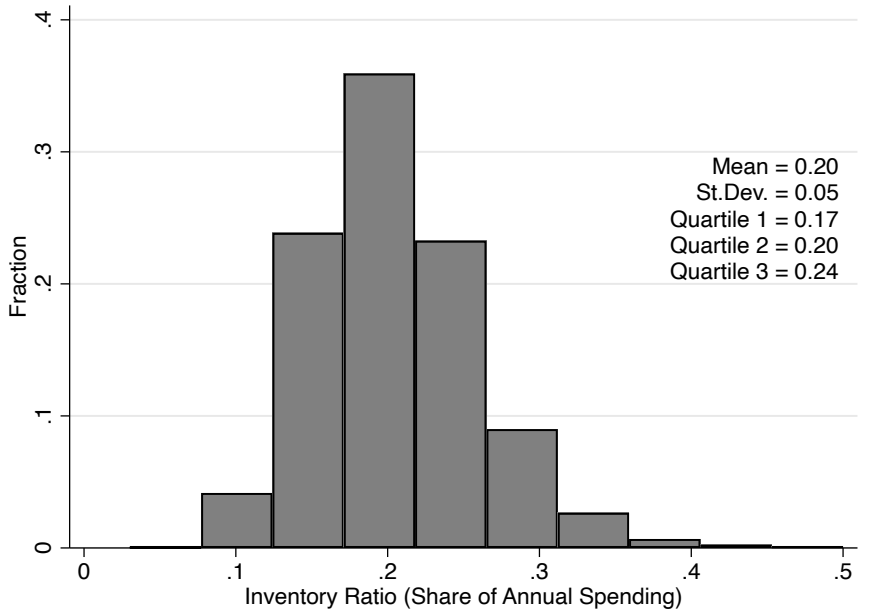
Zinman, Jonathan, "Household Debt: Facts, Puzzles, Theories, and Policies," *Annual Review of Economics*, 2015, *7* (1), 251–276.

FIGURE I
OBSERVED CONSUMER GOODS INVENTORY

(a) Distribution of Inventory Levels



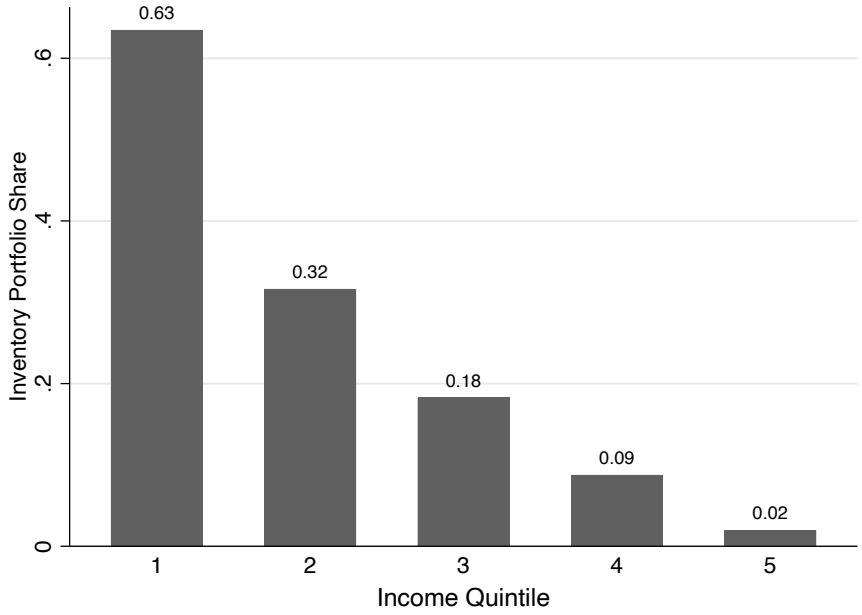
(b) Distribution of Inventory-to-Spending Ratio



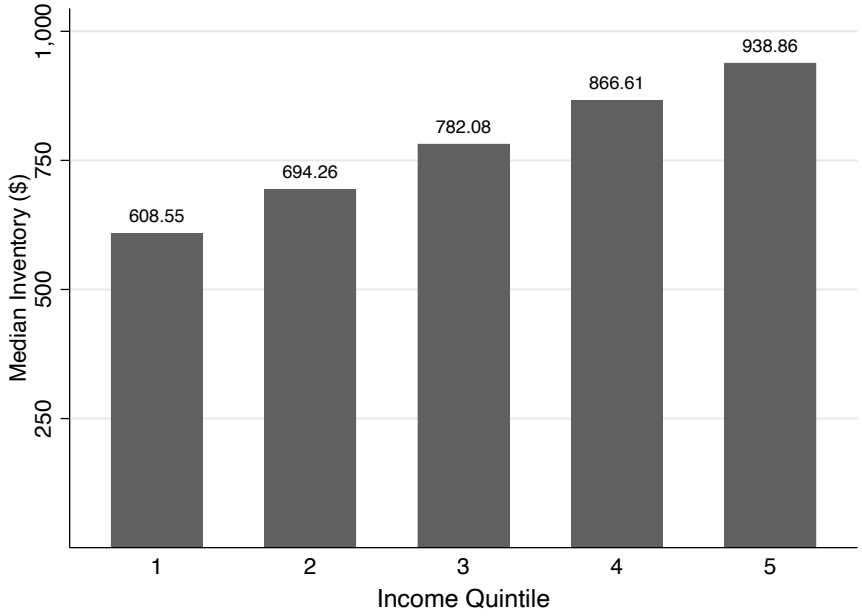
Notes: We compute the average value of inventory for households in the NCP over 2013 and 2014. Panel (a) plots how the average inventory level varies across households. Average inventory is plotted up to the 99th percentile. Summary statistics are computed using the raw data. Panel (b) plots the distribution 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



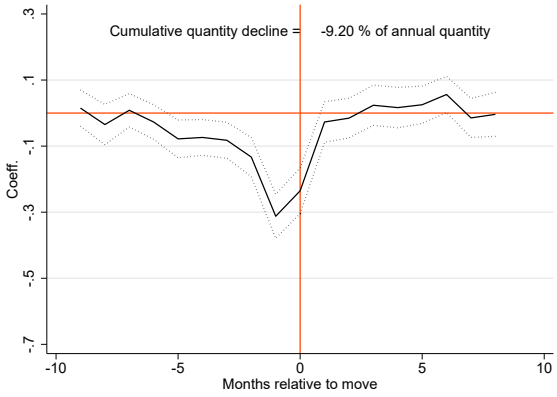
(b) Median Inventory



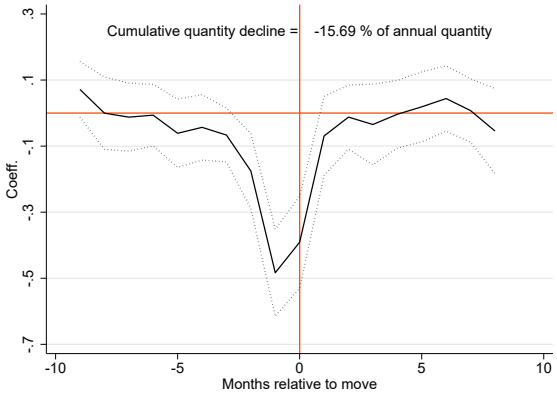
Notes: This figure is constructed by combining data from the NCP over 2013 and 2014 and the SCF over 2010, 2013 and 2016. We compute household income quintiles using the SCF and use household income reported by the Nielsen panelists to assign them to a quintile. We then compute the average value of inventory 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) $\text{Financial Assets}_q$, and also the corresponding inventory portfolio share, $\text{Inventory}_q/(\text{Financial Assets}_q + \text{Inventory}_q)$.

FIGURE III
VALIDATION: LOG QUANTITY PURCHASED AROUND MOVE DATES

(a) All households moving to new 3-digit ZIP Code



(b) Households in top quartile of move distance

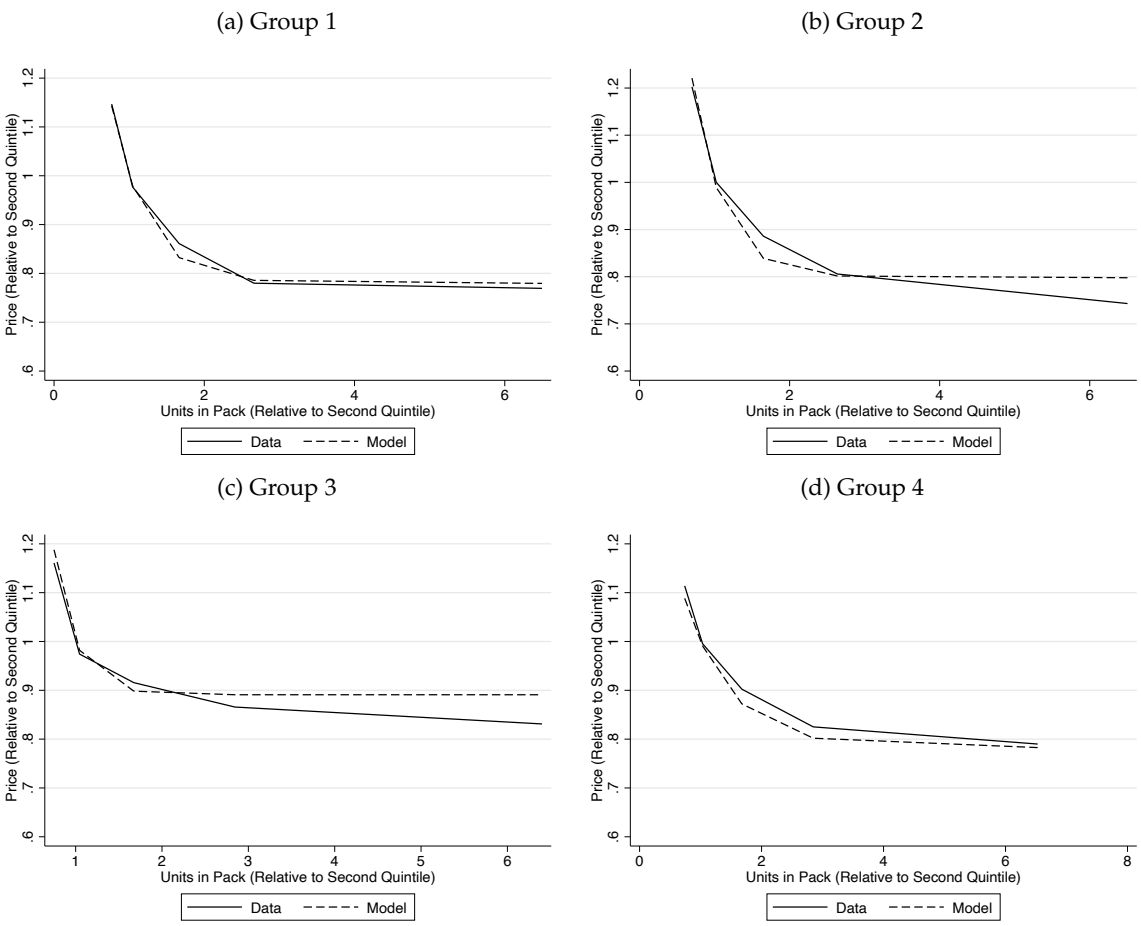


Notes: This figure shows the change in log quantity purchased around the time a household moves. For households who move to a new 3-digit ZIP Code in a given year we impute the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather than their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

$$\log q_{i,t} = \sum_{s=-9}^{9} \beta_s \, \text{Moved}_{i,t-s} + \text{Month FE} + \text{Household FE} + \epsilon_{i,t},$$

where $\log q_{i,t}$ is 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. The sample includes non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. Panel (b) includes only households in the top quartile of move distance – that is households moving more than 975km. Standard errors are clustered by household. Nielsen sampling weights are used.

FIGURE IV
BULK CALIBRATION



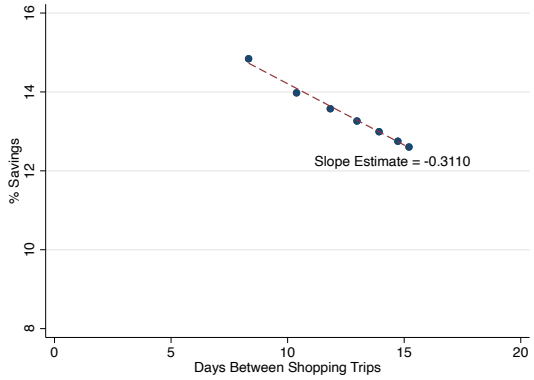
Notes: We first compute pack size quintiles for each product, and then compute the median number of units and unit price for each product and pack size quintile, weighted by UPC expenditure. We normalize both prices and units by dividing by the second quintile price and units. The range of available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

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

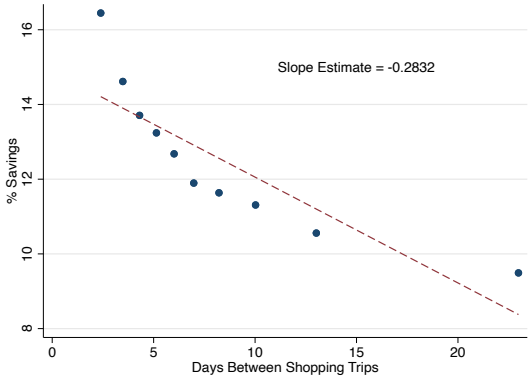
Where p is the product ID described in Section 4.5 and q is the pack size bin. We weight each observation by total spending on product p. The dashed line shows the relative price we assume in the model: Price $= \alpha + \beta e^{-\hat{\sigma}\text{Units}}$ and Units is the weighted average normalized units in units bin q across all products. The solid line is constructed by computing the weighted average normalized retail price in units bin q across all products.

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

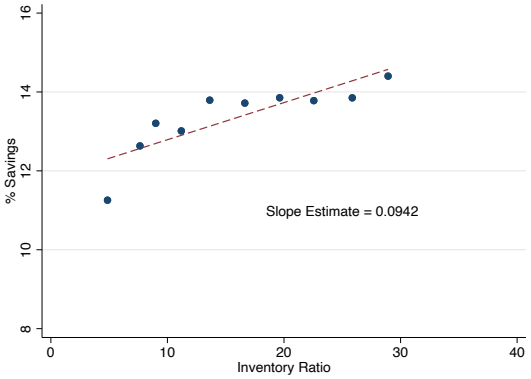
(a) Model savings & trip interval



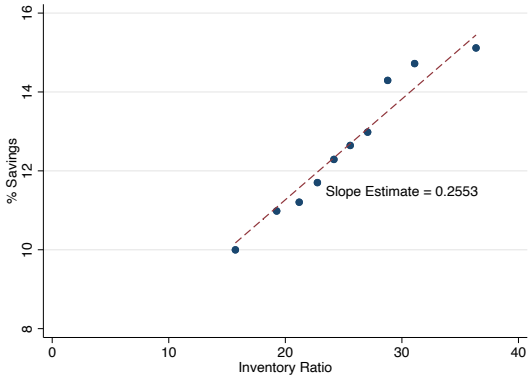
(b) Data savings & trip interval



(c) Model savings & working capital



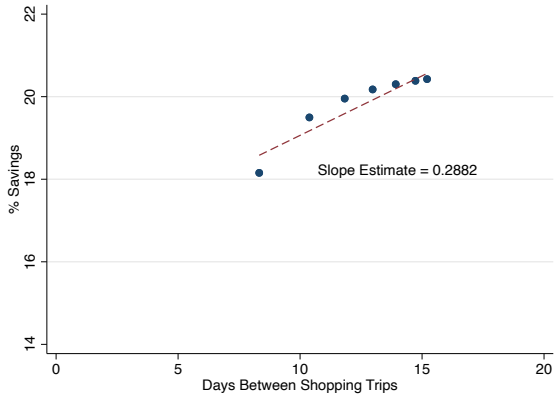
(d) Data savings & working capital



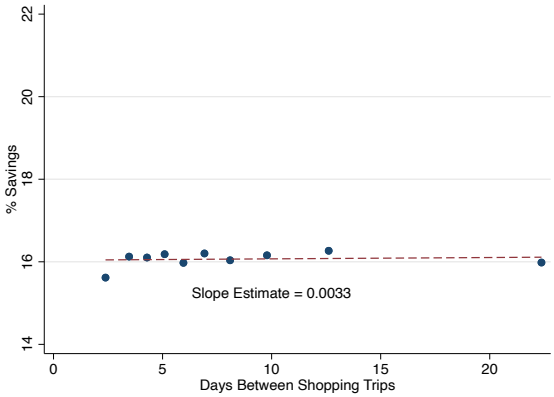
Notes: In Panel (a) we evaluate % deal savings and days between trips in the model for different values of k. Working capital is set sufficiently high that the working capital constraint does not bind. Days between trips is computed as $\Delta \times \frac{365}{12}$. In Panel (b) we plot the data relationship 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 model is average inventory divided by annual spending. 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 30. 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

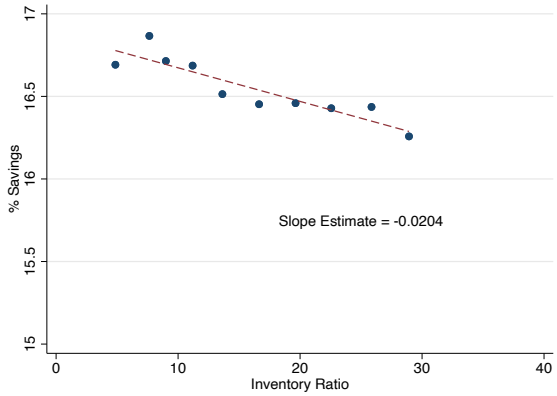
(a) Model savings & trip interval



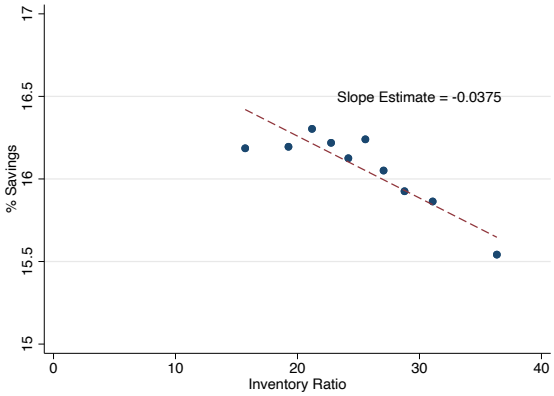
(b) Data savings & trip interval



(c) Model savings & working capital

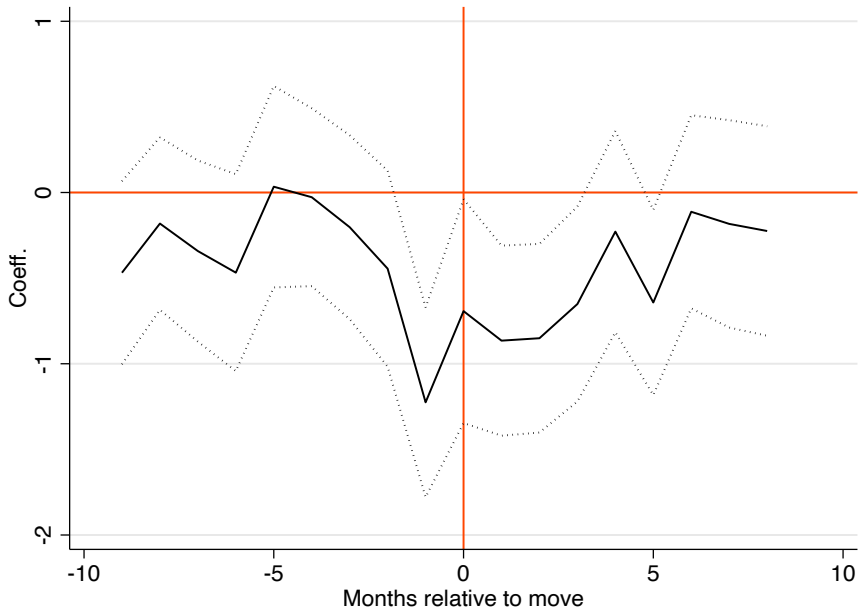


(d) Data savings & working capital



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 model is average inventory divided by annual spending. 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 (31). 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

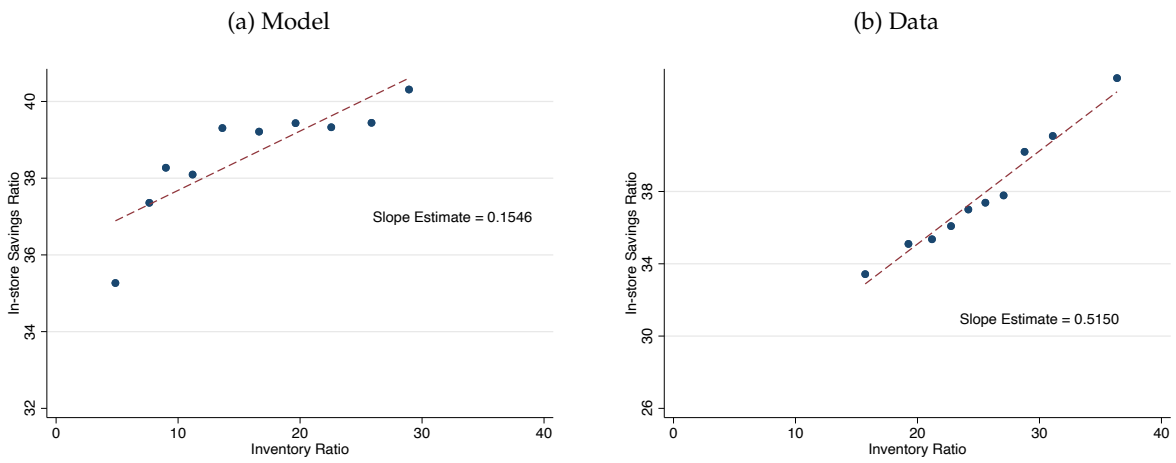


Notes: This figure shows the percentage point change in deal savings around the time a household moves. For households who move to a new 3-digit ZIP Code in a given year we identify the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather 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:

$$\text{Deal Savings}_{h,t} = \sum_{s=-9}^{9} \beta_s \, \text{Moved}_{h,t-s} + \text{Month FE} + \text{Household FE} + \epsilon_{h,t},$$

where $\text{Deal Savings}_{h,t}$ is the percentage deal savings received on transactions by household h in month t, winsorized at the first and 99th percentiles. We compute dollar deal savings at the transaction level and then aggregate to household by department by month. The sample includes non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. 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)	(5)
Shelf Life (Months)	0.112***	0.430***	0.418***	0.414***	
	(0.001)	(0.004)	(0.003)	(0.003)	
Shelf Life Squared		-0.005***	-0.005***	-0.005***	
		(0.000)	(0.000)	(0.000)	
Number of Trips				-1.936***	
				(0.078)	
Shelf Life > 6 months					6.026***
					(0.033)
Household FE			X		
Number of Observations	4,651,322	4,651,322	4,651,321	4,651,322	4,651,322
Adjusted R-squared	0.06	0.08	0.26	0.13	0.06

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

$$\text{Inventory Ratio}_{h,g} = \beta_1 \text{Shelf Life}_g + \beta_3 X_{h,g} + \epsilon_{h,g}.$$

$\text{Inventory Ratio}_{h,g}$ is the ratio of household inventory to annual spending in product group g, multiplied by 100. Column 3 shows results with household fixed effects. Standard errors are clustered by household. Regressions are weighted. The weights are Nielsen sampling weights multiplied by total product group expenditures. * $p < .1$, ** $p < .05$, *** $p < .01$.

TABLE II
FACTORS CORRELATED WITH HOUSEHOLD INVENTORY TO SPENDING RATIOS

VARIABLES	(1) Inv. Ratio	(2) Inv. Ratio	(3) Inv. Ratio	(4) Inv. Ratio
Age	-0.00432***	-0.00432***	-0.00430***	-0.00431***
	(9.48e-05)	(9.48e-05)	(9.47e-05)	(9.48e-05)
Children	0.00650***	0.00652***	0.00636***	0.00638***
	(0.000722)	(0.000722)	(0.000722)	(0.000722)
Married	-0.0140***	-0.0140***	-0.0140***	-0.0140***
	(0.000380)	(0.000381)	(0.000380)	(0.000380)
Dual-Career	0.00134***	0.00136***	0.00137***	0.00138***
	(0.000446)	(0.000446)	(0.000446)	(0.000446)
Home Prices	4.65e-06***	4.71e-06***	4.42e-06***	4.47e-06***
	(7.16e-07)	(7.17e-07)	(7.17e-07)	(7.17e-07)
Pop Density	2.94e-07***	2.96e-07***	3.05e-07***	3.06e-07***
	(3.49e-08)	(3.49e-08)	(3.49e-08)	(3.49e-08)
Single Family Home	0.00583***	0.00584***	0.00579***	0.00580***
	(0.000384)	(0.000384)	(0.000384)	(0.000384)
Income	7.57e-08***	7.49e-08***	7.39e-08***	7.32e-08***
	(5.22e-09)	(5.24e-09)	(5.23e-09)	(5.24e-09)
White	-0.0116***	-0.0116***	-0.0116***	-0.0115***
	(0.000399)	(0.000399)	(0.000398)	(0.000399)
Asian	0.0120***	0.0120***	0.0119***	0.0119***
	(0.000916)	(0.000916)	(0.000915)	(0.000915)
Household Size	-0.00322***	-0.00318***	-0.00314***	-0.00311***
	(0.000148)	(0.000150)	(0.000149)	(0.000150)
Average Product Storability		4.48e-05**		3.80e-05*
		(2.14e-05)		(2.14e-05)
Potential Bulk Savings			-0.0144***	-0.0143***
			(0.00215)	(0.00215)
Observations	115,533	115,533	115,492	115,492
R^2	0.064	0.065	0.065	0.065
Year FE	YES	YES	YES	YES

Notes: Age is the maximum age of household heads. Young children is an indicator for whether children under the age of 6 are present in the house. Dual career is an indicator for both household heads being in the labor force and employed. HPI represents a zipcode level home price index from Redfin. Density represents city-level population density. Single family home is an indicator for whether the household lives in a single family home. Income is binned according to Nielsen standards. White and Asian are indicator variables for the household head. Household size is a count of the number of household members living in the house. Average Product Storability notes the average storability and/or durability of all products purchased by the household over the year. Potential Bulk Savings denotes the amount of money possible to be saved by the household by purchasing in bulk, conditional on the actual product purchases that the household engaged in throughout the year. * $p < .1$, ** $p < .05$, *** $p < .01$.

TABLE III
DEPRECIATION PARAMETERS

Perishability (l)	δ_l	Shelf Life (Months)	$\bar{t}_l$	s_l
1	4.26	0.15	0.25	7.93
2	1.77	0.39	0.39	9.59
3	0.31	2.73	2.73	11.48
4	0.03	36.41	36.41	71.01

Notes: This table shows the model depreciation parameters by perishability group. Group 1 contains the most perishable products and Group 4 contains the least perishable. The calibration approach is described in Section 4.5.2. s_l is the % of total Nielsen spending accounted for by each group. In the model, one unit of a product from group l stored for a period of t months since purchase provides consumption of:

$$\begin{cases} e^{-\delta_l t} \text{ units} & \text{if } t < \bar{t}_l \\ 0 \text{ units} & \text{if } t \geq \bar{t}_l \end{cases}$$

TABLE IV
PRICE PARAMETERS

Perishability (l)	x	p^d	p^f
1	0.20	1.05	0.78
2	0.20	1.06	0.78
3	0.26	1.07	0.80
4	0.24	1.07	0.78

Notes: This table shows the model price parameters by perishability group. Group 1 contains the most perishable products and Group 4 contains the least perishable. The calibration approach is described in Section 4.5.3. x is the probability of a sale, p^d is the price in the event of a sale and p^f is the full price.

TABLE V
BULK SAVINGS PARAMETERS

Perishability (l)	σ	$\hat{a}_0$	$\hat{a}_1$	α	β
1	2.15	0.78	1.89	0.78	1.89
2	2.45	0.79	2.33	0.80	2.34
3	4.00	0.87	5.80	0.89	5.96
4	1.31	0.80	0.82	0.78	0.81

Notes: This table shows the model bulk savings parameters by perishability group. Group 1 contains the most perishable products and Group 4 contains the least perishable. The calibration approach is discussed in Section 4.5.4.Columns 2-3 show estimates of a_0 and a_1 from Equation (23). σ is chosen to maximize the within R-squared of (23). In the model, we normalize the price of the standard pack size to one, and the price of other pack sizes reflect percentage deviations from the standard pack size. We therefore set $\alpha = \frac{\hat{a}_0}{\hat{a}_0+\hat{a}_1 e^{-\hat{\sigma}}}$ and $\beta = \frac{\hat{a}_1}{\hat{a}_0+\hat{a}_1 e^{-\hat{\gamma}}}$. The calibration procedure is discussed in more detail in Section 4.5.4.

TABLE VI
FINANCIAL RETURNS TO HOUSEHOLD INVENTORY INVESTMENT

Working Capital (% of Annual C)	Min. Inv. (%)	% Savings: Deal	Bulk	Interval Δ^* (Months)	m_4^*	Net Return (%)
2.5	45.59	8.57	14.30	0.20	3	172.32
5.0	71.27	12.43	15.85	0.22	6	59.41
7.5	78.96	13.46	17.12	0.24	8	25.64
10.0	83.61	14.06	17.54	0.25	10	15.60
12.5	87.83	14.68	16.89	0.24	13	8.07
15.0	89.63	14.82	17.14	0.24	15	4.38
17.5	90.98	14.89	17.30	0.25	17	2.38
20.0	92.02	14.93	17.42	0.25	19	1.29
22.5	92.84	14.95	17.50	0.25	21	0.71
25.0	93.51	14.97	17.51	0.25	23	0.41
27.5	94.07	14.98	17.51	0.25	25	0.24
30.0	94.54	14.99	17.51	0.25	27	0.14
32.5	94.94	14.99	17.51	0.25	29	0.08
35.0	95.28	14.99	17.51	0.25	31	0.05

Notes: This table is constructed by solving the model for different values of $\bar{I}$. The working capital ratio in column 1 is working capital in ($\bar{I}$) expressed as a percentage of annual consumption. Column 2 shows the value of inventory immediately prior to a trip as a percentage of inventory plus cash required to cover the total price paid in store in \$, $\sum_l P_l \cdot S_l$. 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_4^* is the optimal deal shopping strategy for goods with a shelf life of at least seven months. The net return incorporates not only in-store savings but also depreciation and trip fixed costs.

TABLE VII
FINANCIAL RETURNS WITH HIGH SHOPPING TRIP FIXED COST ($k = 0.1$)

Working Capital (% of Annual C)	Min. Inv. (%)	% Savings: Deal	% Savings: Bulk	Interval Δ^* (Months)	m_4^*	Net Return (%)
2.5	13.13	3.26	18.81	0.30	1	344.99
5.0	44.75	7.83	19.92	0.38	3	92.43
7.5	63.73	10.24	19.94	0.39	5	41.22
10.0	69.32	10.81	20.15	0.42	6	36.27
12.5	77.64	12.00	20.04	0.40	8	16.42
15.0	80.19	12.21	20.14	0.42	9	12.56
17.5	84.38	12.85	20.02	0.40	11	10.18
20.0	85.72	12.89	20.10	0.41	12	5.01
22.5	86.89	12.95	20.14	0.42	13	3.36
25.0	87.93	13.04	20.14	0.42	14	3.95
27.5	89.65	13.22	20.10	0.41	16	2.05
30.0	90.27	13.20	20.14	0.42	17	1.13
32.5	90.87	13.23	20.14	0.42	18	0.86
35.0	91.39	13.25	20.14	0.42	19	1.12

Notes: See the description in Table VI.

Online Appendix

Financial Returns to Household Inventory Management

A Savings Measures

This section describes how we compute deal and bulk savings using the NCP and NRP. We begin by computing 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 that 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:

$$\text{AvgSecondQuintilePrice}_{z,p,y} = \frac{1}{|U_p|} \sum_{u \in U_p} \text{AvgPrice}_{z,u,y} \tag{28}$$

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 4.5. $\text{AvgPrice}_{z,u,y}$ is the average price at which UPC u was sold in year y in 3-digit Zip Code z:

$$\text{AvgPrice}_{z,u,y} = \frac{\sum_{w \in W_y} \sum_{s \in S_{z,w}} P_{u,w,s} 1[Q_{u,w,s} > 0]}{\sum_{w \in W_y} \sum_{s \in S_{z,w}} 1[Q_{u,w,s} > 0]} \tag{29}$$

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 is 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 ($Q_{u,w,s} = 0$).

A.1 Deal Savings

We compute annual deal savings using the following formula:

$$\%\text{Deal savings}_{h,y} = \frac{\sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t} - \text{Spending}_{h,y}}{\sum_{t \in T_{h,y}} \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}} \tag{30}$$

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.[31] $Q_{u,h,t}$ is the quantity of

[31] Using 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

UPC u purchased by household h on trip t and $\sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}$ 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 behavior 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.

A.2 Bulk Savings

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

$$\%\text{Bulk savings}_{h,y} = \frac{\sum_t \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t} - \sum_t \sum_u \text{AvgPrice}_{z,u,y} \cdot Q_{u,h,t}}{\sum_{t \in T_{h,y}} \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t}} \tag{31}$$

That is, the bulk savings measure compares the average price of the UPC actually purchased with the average price of a second quintile pack size of the same product.

A.3 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:

$$\%\text{Total savings}_{h,y} = \frac{\sum_{t \in T_{h,y}} \sum_p \text{AvgSecondQuintilePrice}_{z,p,y} \cdot Q_{p,h,t} - \text{Spending}_{h,y}}{\text{Spending}_{h,y}} \tag{32}$$

B 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 con-

value of coupon discounts.

sumption 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 C.6a, 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 C.6b 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 C.6a. 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 C.7a, 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 C.7b, 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 C.8a 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).

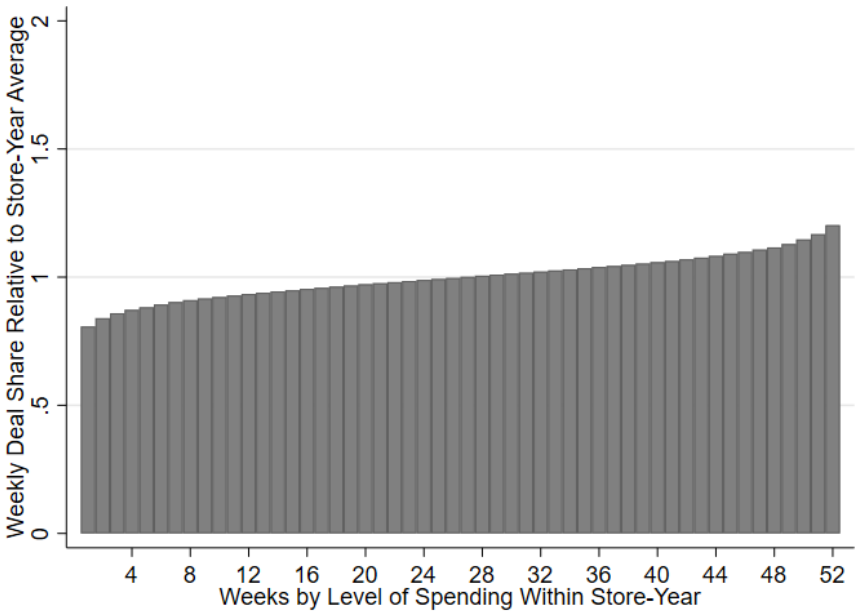
We also separate respondents where one interview day was a weekday and the other day was a weekend. Even without further aggregation the share of items also consumed on Day 2 is quite high, at around 40%. The share declines slightly with the time between interviews, consistent with some of the persistence being driven by households consuming items from the same shopping trip, but remains high even with a gap of 9 days. Figure C.8c 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 C.8e) this further increases to 90%.

NHANES respondents also report the amount of each item consumed in grams. Figures C.8b, C.8d and C.8f 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).

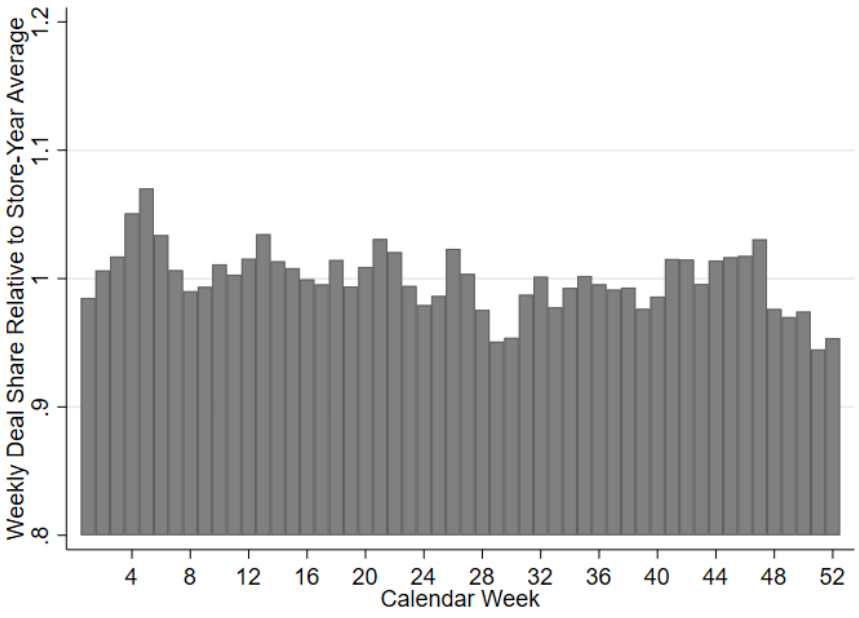
C Additional Appendix Figures and Tables

FIGURE C.1
RETAILER DEAL CONCENTRATION

(a) Ranked Weeks



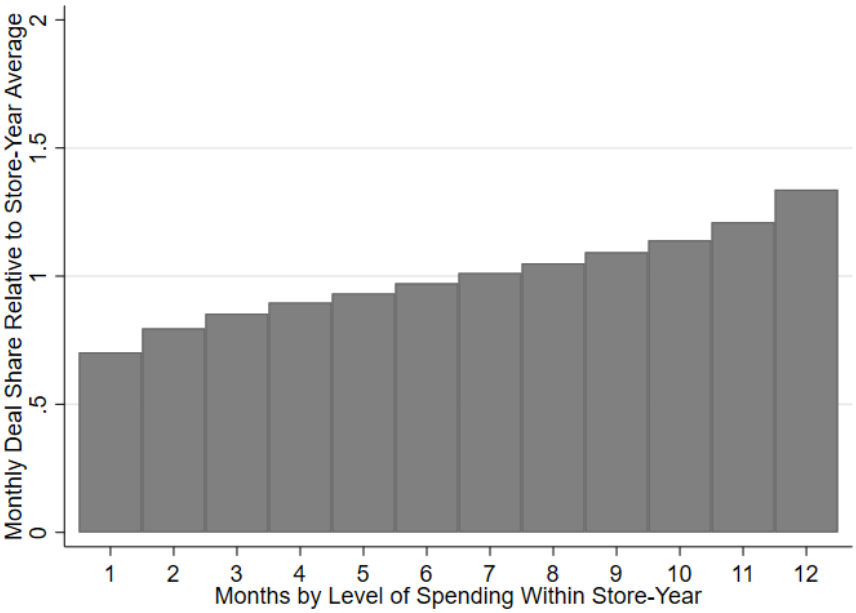
(b) Calendar Weeks



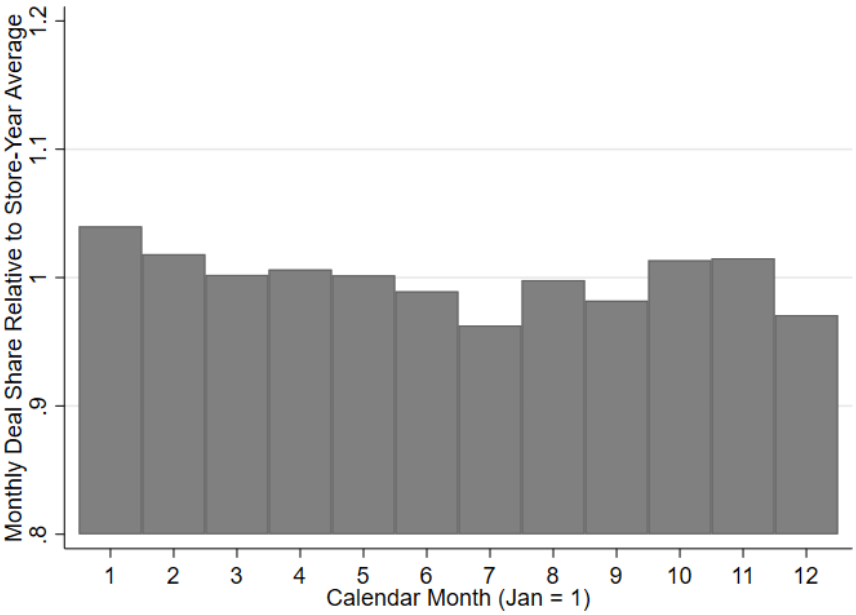
Notes: We compute the share of deal sales for each retailer in each week using the deal flag in the NCP (which includes both coupon and non-coupon deals), and then divide by the retailer's average deal share over the year. Panel (a) plots the average across retailers by ranked weeks (so week 1 is the 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 C.2
STORE DEAL CONCENTRATION

(a) Ranked Months



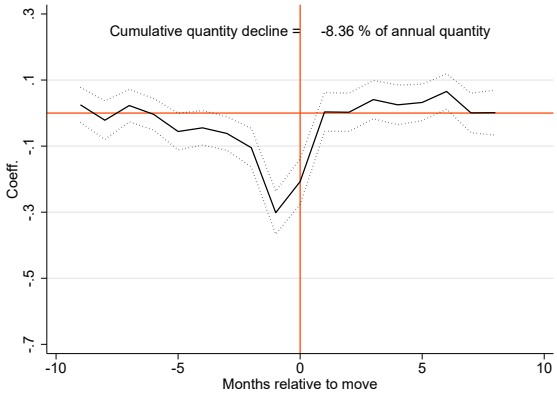
(b) Calendar Months



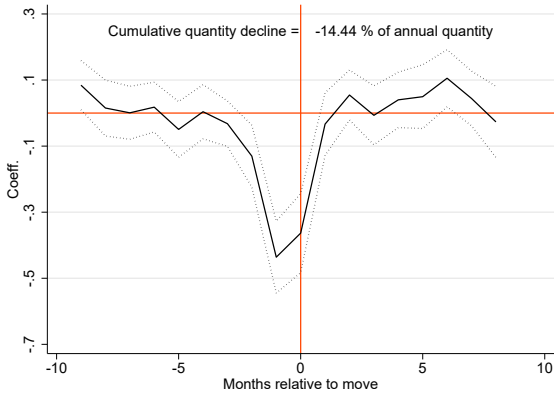
Notes: We compute the share of deal sales for each store in each month using the deal flag in the NCP (which includes both coupon and non-coupon deals), and then divide by the store's average deal share over the year. Panel (a) plots the average across stores by ranked month (so month 1 is the month 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 C.3
VALIDATION: LOG QUANTITY PURCHASED AROUND MOVE DATES

(a) All households moving to new 3-digit ZIP Code

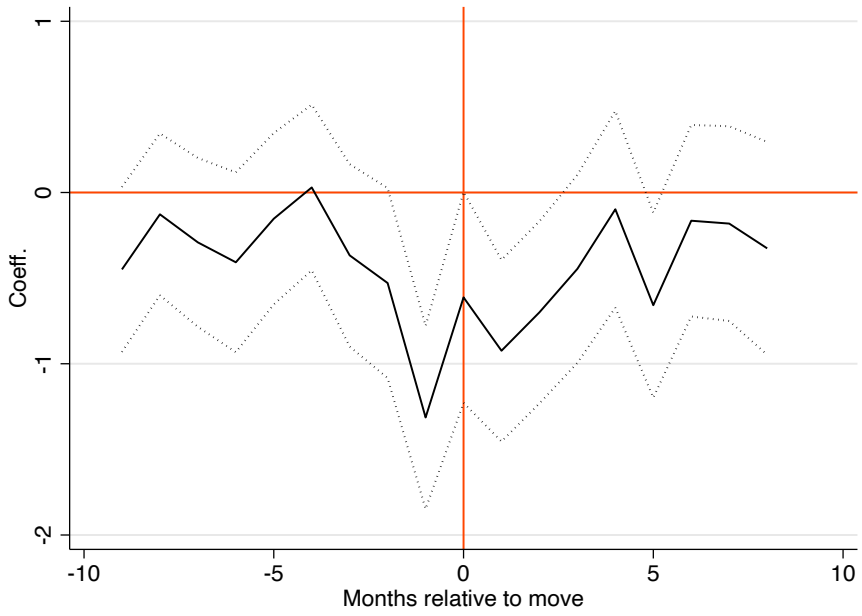


(b) Households in top quartile of move distance



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 responses analogous to Figure III estimated using the imputation approach described by Borusyak et al. (2021). We implement this using their Stata command, did_imputation. The sample includes non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. Panel (b) includes only households in the top quartile of move distance – that is households moving more than 975km. Standard errors are clustered by household. Nielsen sampling weights are used.

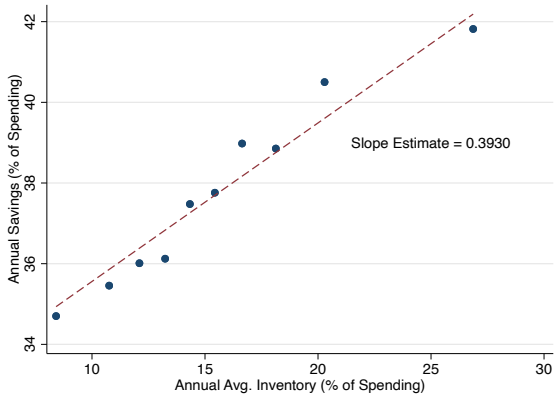
FIGURE C.4
DEAL SAVINGS AROUND A MOVE



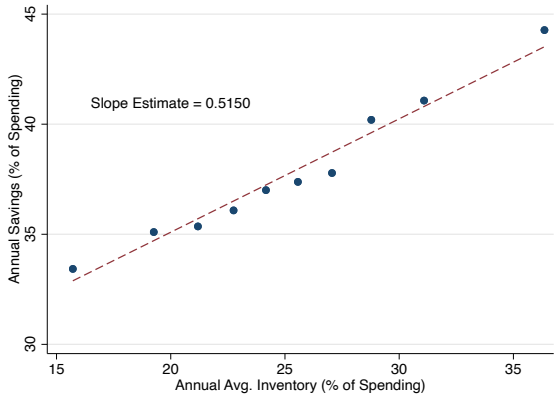
Notes: This figure shows the percentage point change in deal savings around the time a household moves. For households who move to a new 3-digit ZIP Code in a given year we identify the month of the move by searching for a break in the share of trips made in the household's new 3-digit ZIP Code (rather than their old 3-digit ZIP Code). The figure plots responses analogous to Figure VII estimated using the imputation approach described by Borusyak et al. (2021). We implement this using their Stata command, did_imputation. We compute dollar deal savings at the transaction level and then aggregate to household by department by month. The sample includes non-movers and households who moved to a new 3-digit ZIP Code exactly once between 2006 and 2014. The sample period is January 2006 to December 2014. We also drop households who leave the panel and re-enter in a later year. Standard errors are clustered by household. The regression is weighted using Nielsen sampling weights.

FIGURE C.5
GROSS RETURN UNDER ALTERNATIVE AGGREGATION ASSUMPTIONS

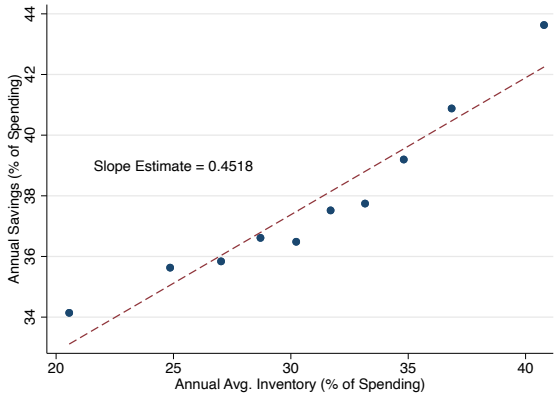
(a) Nielsen Department



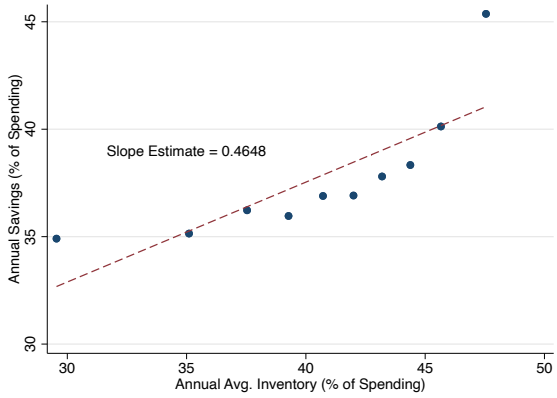
(b) Nielsen Product Group



(c) Nielsen Product Module



(d) Product Code (UPC)



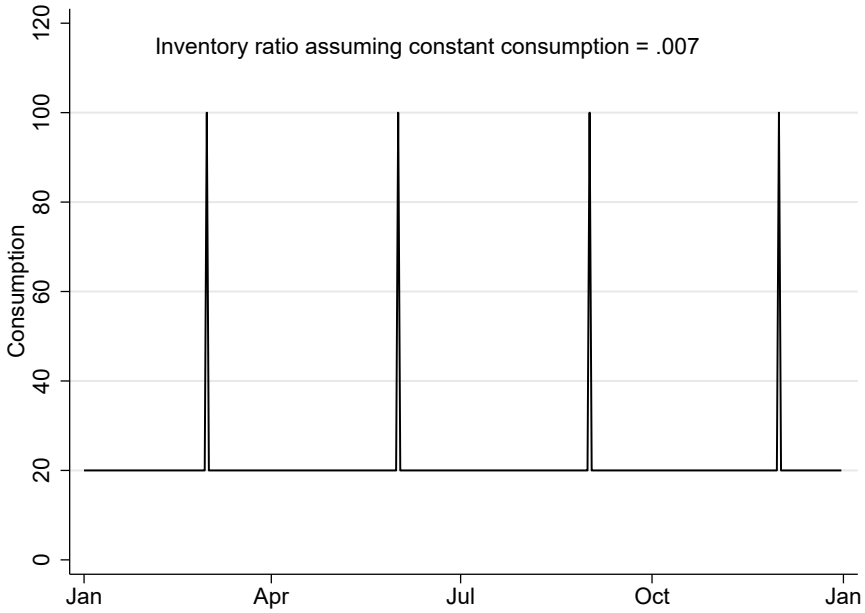
Notes: This figure uses the NCP over 2013 and 2014 to illustrate the relationship between in-store savings and average inventory as a percentage of spending. Average inventory is computed with decreasing level of aggregation, starting from the 10 "Nielsen Departments" and going to 118 "Nielsen Product 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 depreciation costs or trip fixed costs. The red dotted line shows predicted values from the following regression specification:

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

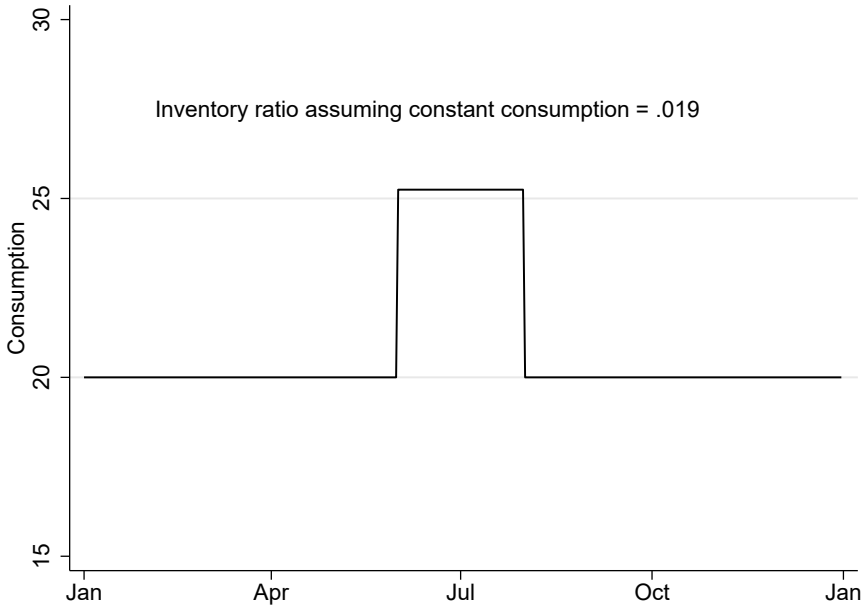
Regressions are weighted using Nielsen sampling weights.

FIGURE C.6
NON-CONSTANT AGGREGATE CONSUMPTION

(a) Example 1



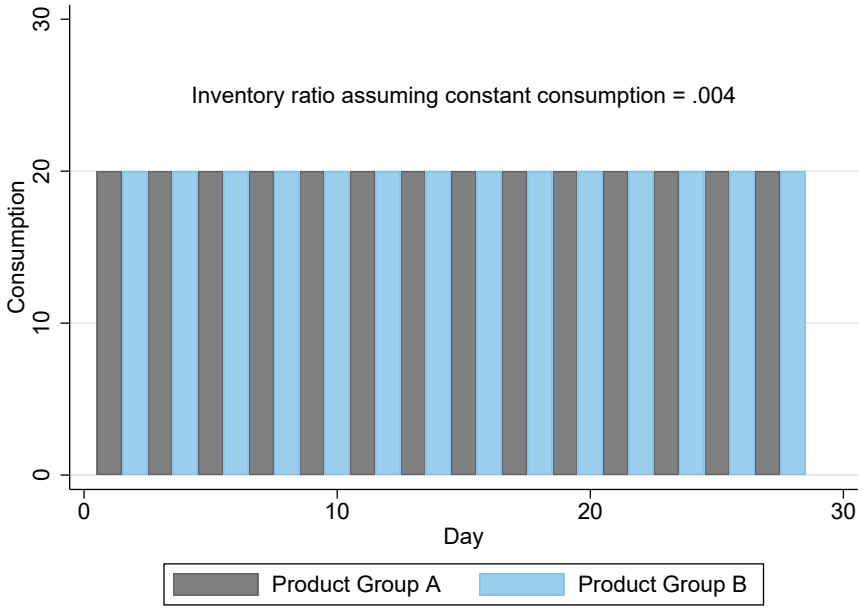
(b) Example 2



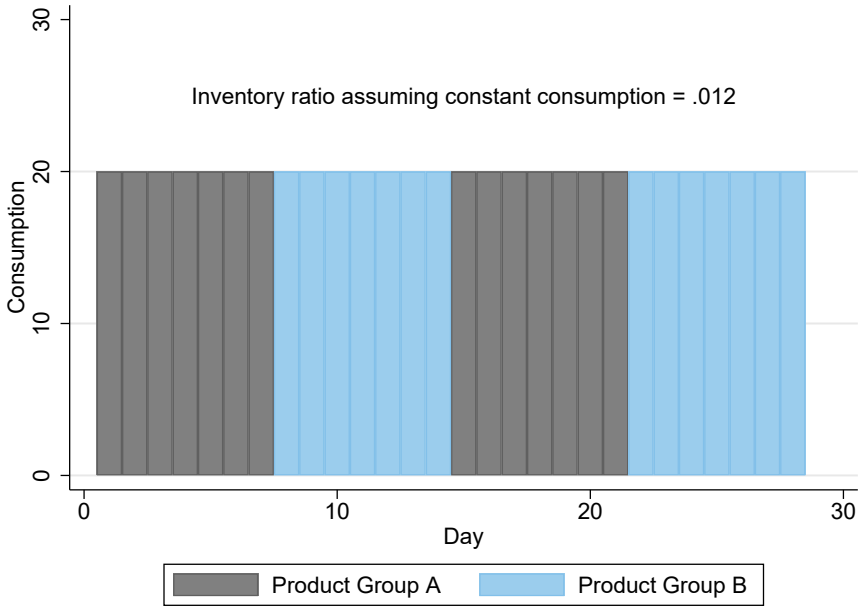
Notes: Panel (a) plots consumption for a household who consumes five times more on the first day of March, June, September and December than on other days of the year. Panel (b) plots consumption for a household who consumes around 25 per cent more in June and July than in other months. The inventory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

FIGURE C.7
PRODUCT GROUP SWITCHING

(a) Example 1



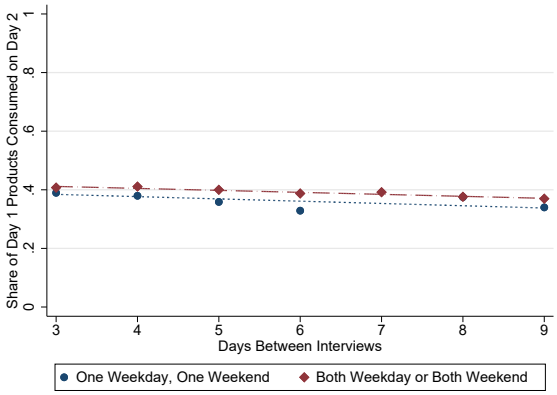
(b) Example 2



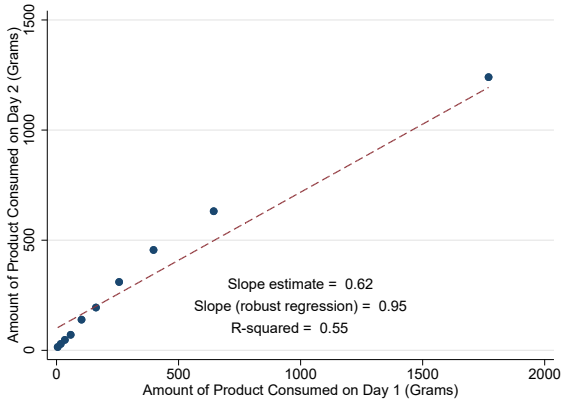
Notes: Panel (a) shows the consumption pattern of a household who consumes only Product Group A on one day and only Product Group B on the following day. Panel (b) shows the consumption patter of a household who consumes only Product Group A one week and only Product Group B the following week. Aggregate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

FIGURE C.8
VALIDATION: CONSUMPTION PERSISTENCE

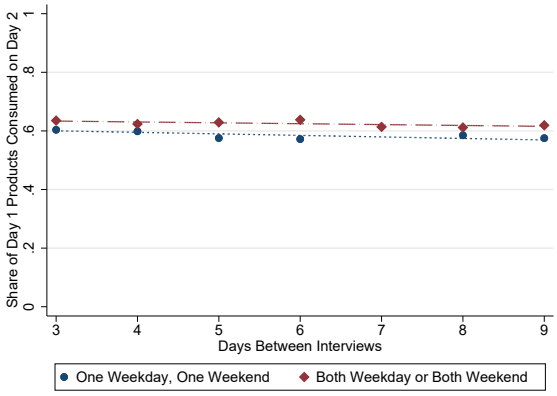
(a) NHANES Product Code



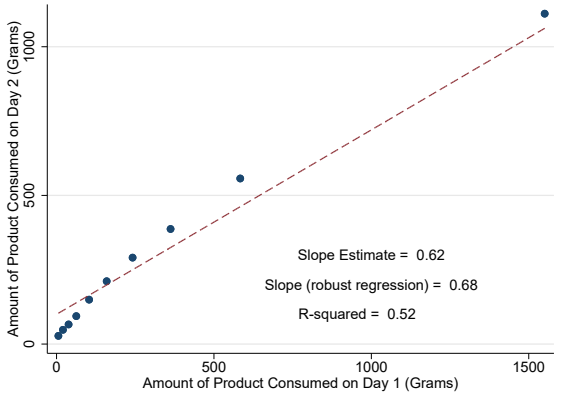
(b) NHANES Product Code



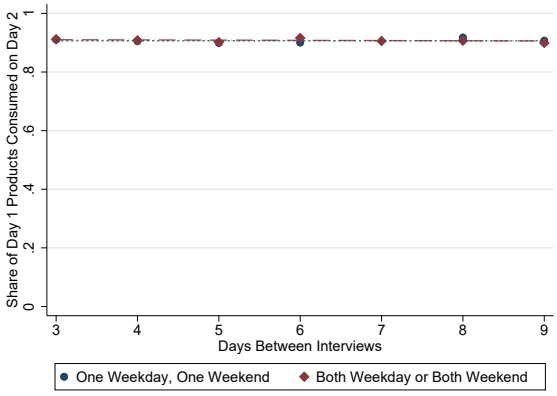
(c) Nielsen Product Group



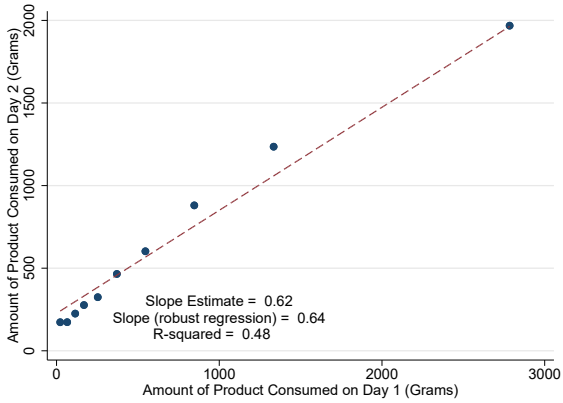
(d) Nielsen Product Group



(e) Nielsen Department



(f) Nielsen Department



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 C.1
HOME PRICES AND RESTAURANT SPENDING

	(1)	(2)	(3)	(4)
	Share of Household Consumer Spending at Restaurants			
ln(House Price - City)	0.0360***	0.0367***	0.0365***	0.0354***
	(0.00137)	(0.00135)	(0.00135)	(0.00272)
ln(Annual Income)		-0.0320***	-0.0322***	-0.0331***
		(0.00116)	(0.00116)	(0.00114)
Observations	28,212	28,212	28,212	28,212
R^2	0.024	0.049	0.050	0.098
Year FE	NO	NO	YES	YES
State FE	NO	NO	NO	YES

Notes: * $p < .1$, ** $p < .05$, *** $p < .01$.

TABLE C.2
PRODUCT DEPRECIATION

Product	Quality changes over time
Refrigerated prepared coleslaw. *Nielsen group:* dressings, salads, prep foods-deli. *Group FSIS shelf life:* 5 days.	Several quality outcomes change rapidly from date of purchase (increase in acidity, decline in dressing thickness and viscosity, increase in cabbage translucency). Flavor is broadly unchanged for 4-6 days and then starts to decline. Despite quality deterioration, is safe to eat well past expiry date if stored at 5 degrees celsius (Brocklehurst, 1994).
Sourdough Bread *Nielsen group:* bread and baked goods. *Group FSIS shelf life:* 5 days.	Perceived quality starts to deteriorate immediately from date of purchase (less crisp, more difficult to chew and swallow, crumb dries out). The share of consumers who would not eat the bread increases from around 20% at 1.5 days to around 50% at 3 days and around 90% at 5 days (Gauchez, Loiseau, Schlich and Martin, 2020).
Liquid milk *Nielsen group:* milk. *Group FSIS shelf life:* 10 days.	In lab conditions milk stored at home refrigeration temperatures has a very low rejection rate in the first 5 days. Rejection increased to 100% by 17.5 days. However, as milk deteriorates much faster at higher temperatures (Duyvesteyn, Shimoni and Labuza, 2001), in practice there could be considerable variation in shelf life due to variation in storage both before and after purchase (Lewis and Hale, 1994). For example, consumers may store milk in the refrigerator door, which tends to be warmer (Terpstra, Steenbekkers, de Maertelaere and Nijhuis, 2005).
Frozen grass-fed lamb *Nielsen group:* Frozen unprepared meat/poulty/seafrood *Group FSIS shelf life:* 9 months	No change in quality for at least two years (Winger, 1984). In general, changes in quality for frozen food are slow and food going off is not the primary concern. Shelf-life should reflect the time period over which consumers perceive the product to have the expected level of quality (Symons, 1994).
Breakfast cereal *Nielsen group:* cereal. *Group FSIS shelf life:* 9 months.	No sign of off flavors from 0-14 months. No change in appearance, texture or flavor over the first 4 months. Possibly some, but not substantial, deterioration in these outcomes between 4 and 14 months. Similar for cereals containing fruit, except the moisture of the fruit starts to decline immediately (though flavor scores were stable over 14 months) (Howarth, 1994).
Fruit-filled snack bar *Nielsen group:* breakfast food. *Group FSIS shelf life:* 11.25 months.	No decline in acceptability before 35 weeks when stored at 20 degrees celsius. Faster decline when stored at 30 degrees celsius (Corrigan, Hedderley and Harvey, 2012)
Ready-meal (pasta in minced meat and tomato sauce) *Nielsen group:* prepared food ready-to-serve. *Group FSIS shelf life:* 20 months.	The product is sterilized so changes in texture and flavor determine expiry. Acidity from the tomato contributes to faster deterioration than for similar products. No decline in acceptability scores for the first 3 months when stored at 25 degrees celsius (6 months if refrigerated). Acceptability then declined very slowly up to the end of the trial at 18 months. The product still had high acceptability at 18 months when stored at 25 degrees celsius or less (Goddard, 1994).