

Financial Returns to Household Inventory Management

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

August 2020

Abstract

Households tend to hold substantial amounts of non-financial assets in the form of inventory. Households can obtain significant financial returns from strategic shopping and optimally managing these inventories of consumer goods. In addition, they choose to maintain liquid savings – household working capital – not just for precautionary motives but also to support this inventory management. We demonstrate that households earn high returns from inventory management at low levels of inventory, though returns decline rapidly as inventory levels increase. This offers a new rationale for poorer households not to participate in risky financial markets, while wealthier households invest in both financial assets and working capital. We provide evidence using scanner and survey data that supports this conclusion.

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

Keywords: household working capital, stock market participation, financial returns, inventory, stockpiling

*Baker: Northwestern University, Department of Finance, 2169 Campus Drive, Evanston, IL 60208; s-baker@kellogg.northwestern.edu. Johnson: 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 Institute (SFI), Northwestern University, NBER, and CPER, Via Giuseppe Buffi 13, CH-6900 Lugano TI, Switzerland; lorenz.kueng@usi.ch. 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 wealth in the form of consumer good inventories. These inventories can be managed over time through strategic shopping behavior as households are able to take advantage of coupons, temporary low prices at retailers, and savings from buying in bulk. Aggregating across all Nielsen Homescan goods, we estimate that households hold approximately \$1,100 in consumer goods inventory at any given time, representing an unmeasured source of non-financial wealth. For households in the lowest quintile of household income, this inventory likely makes up a majority of total household wealth. Moreover, households can earn high returns through the maintenance of liquid savings and strategic shopping behavior.

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

As such, household inventory management offers a new explanation of 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)). This complementary explanation is based on a previously ignored asset class – household working capital – which offers marginal returns that dominate financial market returns at low levels of investment. Importantly, the return on working capital is investor-specific, approximately risk-free, and declines as wealth and inventory holdings increase. It can therefore help explain the strong positive correlation between stock market participation rates and household wealth observed in most countries (e.g. Badarinza, Campbell and Ramadorai (2016)). Furthermore, it can help explain why poorer households choose not to participate in the stock market at all, without resorting to participation fixed costs which are currently the main explanation put forth in the literature (e.g. Gomes, Haliassos and Ramadorai (forthcoming)).

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

Using scanner data from AC Nielsen and income and asset data from the Survey of Consumer Finances (SCF), we provide evidence in support of this new channel. In particular, we compute the total net returns to investment in household working capital. We go one step further than existing work (e.g. Griffith, Leibtag, Leicester and Nevo (2009); Nevo and Wong (2019)), which focuses on in-store savings as a percentage of the product price, but does not take into account the additional

household working capital that must be held to facilitate these savings and the financial returns to this working capital. Moreover, we extend our framework to include the costs from product depreciation and the relation between the level of inventory holdings and differences in shopping trip fixed costs associated with different shopping behaviors.

We build a parsimonious model of inventory management to incorporate these additional components of returns to household working capital investments. The model highlights two sources of returns. By taking larger and less frequent trips, households can save on trip fixed costs and also take advantage of lower unit prices by buying goods in bulk. Alternatively, consumers can shop more frequently, at higher cumulative trip fixed costs, giving them additional opportunities to take advantage of temporary deals at retailers.

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

Existing models of deal shopping focus on individual products in a stochastic framework (e.g. [Boizot, Robin and Visser \(2001\)](#); [Hendel and Nevo \(2013\)](#)). In contrast, we focus on an aggregate deterministic steady state, which is derived from stochastic foundations under the assumption of independent price deals across goods and backed by observations from the data. This has implications for households' cash holdings. 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. Although when focusing on purchases of individual products, fluctuations in prices are observed, when focusing on aggregate shopping trips these fluctuations are smoothed out. In contrast, if aggregated deals are autocorrelated, households may want to hold substantial additional cash to stock up more in those (random) weeks. We provide empirical evidence supporting the former scenario. Stores generally feature consistent amounts of goods on sale throughout the year rather than concentrating deals in particular weeks.¹

The model predicts that households with high shopping fixed costs – such as rich households with a high opportunity cost of time – are more likely to engage in bulk buying and shop less frequently. The proportion of working capital held in liquid financial assets is higher for these households. As the fixed cost of shopping rises from \$1 to \$50 per trip, the cash component of working capital increases by 74%, while the inventory component increases by 40%.

By highlighting the role of household working capital for households' portfolio allocation – which is especially important 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, which has received considerable attention in finance (e.g. [Petersen and Rajan \(1997\)](#); [Fisman and Love \(2003\)](#); [Yang and Birge \(2018\)](#)), inventories of consumers goods and household working capital has been largely ignored by the household finance

¹See Appendix Figures [A.1](#) and [A.2](#).

literature.²

For instance, none of the country studies of household portfolios in the widely cited book by [Jappelli, Guiso and Haliassos \(2002\)](#) include household inventories. This also applies to the chapter by [Bertaut and Starr \(2000\)](#), who study U.S. households' portfolios.³ One explanation for this gap is that inventories are often hard to observe and measure. For example, they are missing from traditional consumer finance data such as the SCF.

Our paper is therefore one of the first systematic studies of the role of household inventories in household finance.⁴ We also highlight how adding household inventory management to a household's portfolio choice problem affects its decision of whether to participate in the stock market. The literature on the participation puzzle is among the oldest in household finance and too large to adequately survey here.⁵ We contribute to this literature by showing that investment in household working capital has returns that vary systematically across households by wealth and that dominate equity returns for poorer and less educated households.⁶

We show that consistent with our model's predictions, observed stock market participation choice of households with higher education is well explained by variation in returns to working capital investments. At low levels of financial assets relative to income, educated households do not participate in the stock market but instead invest largely in liquid assets. As the ratio of financial assets to income increases, the participation rate of these households increases sharply.

Our model implies that participation rates quickly approach 100% as financial wealth increases, while in the data direct participation rates "only" reach slightly above 40%. This is not surprising since our model intentionally abstracts from all other dimensions discussed in footnote 5 that have previously been proposed in the literature to explain non-participation of relatively wealthy households, such as participation costs, correlated background risk, trust, peer effects, etc.

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

³The category "other nonfinancial 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 nonfinancial assets such as artwork, antiques, jewelry, furniture, and valuable collections (coin, stamp, etc.)."

⁴There is a related literature in macroeconomics that studies heterogeneity in the effective unit price paid for similar goods across households and over the business cycle; 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 \(2019\)](#).

⁵The handbook chapters by [Guiso and Sodini \(2013\)](#) and [Beshears, Choi, Laibson and Madrian \(2018\)](#) provide recent surveys of this literature. A non-exhaustive list of explanations of the participation puzzle include pecuniary and non-pecuniary participation fixed costs ([Luttmer \(1999\)](#); [Vissing-Jørgensen \(2002\)](#)); low financial literacy ([Van Rooij, Lusardi and Alessie \(2011\)](#); [Black, Devereux, Lundborg and Majlesi \(2018\)](#)); non-expected utility with first-order risk aversion ([Barberis, Huang and Thaler \(2006\)](#); [Epstein and Schneider \(2010\)](#)); heterogeneity in beliefs ([Kézdi and Willis \(2009\)](#); [Malmendier and Nagel \(2011\)](#); [Hurd, Van Rooij and Winter \(2011\)](#); [Adelino, Schoar and Severino \(2020\)](#)), lack of trust ([Guiso, Sapienza and Zingales \(2008\)](#); [Gennaioli, Shleifer and Vishny \(2015\)](#)), and unawareness of the excess return premium ([Guiso and Jappelli \(2005\)](#); [Grinblatt, Keloharju and Linnainmaa \(2011\)](#); [Cole, Paulson and Shastry \(2014\)](#)); background risk ([Heaton and Lucas \(2000\)](#); [Cocco, Gomes and Maenhout \(2005\)](#)) and positive correlation of stock returns with returns of other assets in household portfolios ([Benzoni, Collin-Dufresne and Goldstein \(2007\)](#); [Davis and Willen \(2014\)](#); [Bonaparte, Korniotis and Kumar \(2014\)](#)); liquidity constraints, illiquid assets and consumption commitments ([Grossman and Laroque \(1990\)](#); [Haliassos and Michaelides \(2003\)](#); [Chetty and Szeidl \(2007\)](#)); or social interactions ([Hong, Kubik and Stein \(2004\)](#); [Kaustia and Knüpfer \(2012\)](#)).

⁶The only other asset class we are aware of that has been suggested to dominate equity returns and to justify non-participation in the stock market – especially early in life – is human capital (e.g. [Roussanov \(2004\)](#); [Athreya, Ionescu and Neelakantan \(2015\)](#)).

In the data, we find that households with low education, such as high school dropouts, do not participate in the stock market and instead maintain a high share of their portfolio invested in liquid assets. Such behavior can be rational in light of previous studies that have shown that households with low education tend to earn low returns on their stock market investments conditional on participation ([Calvet, Campbell and Sodini \(2007\)](#)), for example due to under-diversification, local bias, excessive trading or paying high management fees.

Uneducated household also tend to put low trust in financial markets and financial advisors – often justifiably so (e.g. [Bergstresser, Chalmers and Tufano \(2008\)](#); [Mullainathan, Noeth and Schoar \(2012\)](#); [Anagol, Cole and Sarkar \(2017\)](#); [Linnainmaa, Melzer and Previtero \(forthcoming\)](#)). However, when investing in consumer goods inventory, households do not have to delegate their investment decision. Moreover, they have lots of experience with store price discounts from frequent shopping and are therefore capable of taking advantage of the investment opportunities offered by working capital management. We show that while the returns to working capital investments are similar across households with similar inventory holdings, uneducated households in our data achieve indeed slightly higher in-store savings from shopping strategically.

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

Our finding that working capital investment returns can be very high at low levels of inventory – over 30% per year – suggests that it might be rational for poorer households to temporarily borrow at high interest rates to take advantage of such high returns (e.g. using credit card debt, although probably not payday loans).

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

Alterations in strategic shopping behavior also help explain portions of the “excess sensitivity” of consumption to unanticipated temporary income shocks experienced by households (e.g. [Jappelli and Pistaferri \(2010\)](#)). In the existing consumption literature, many retail purchases, such as grocery store and pharmacy spending, are treated strictly as nondurables. As much of the literature moves to monthly and even daily measures of spending to improve identification of causal effects, the wedge between household spending and actual consumption grows more important. We show that households hold substantial stocks of consumer goods, make purchases in discrete bundles, and

run them down over time. This reinforces the idea that large increases in household spending in one period may translate into increased consumption only over several periods.

The remainder of the paper is structured as follows. Section 2 describes the data sources. Section 3 discusses how we construct our measures of household inventory and the various channels of savings. Section 4 describes our measures of financial returns from investment in inventory. Section 5 lays out the household shopping model and discusses various out of sample predictions. Section 6 uses the model to estimate the financial net returns to investing in household working capital. Section 7 applies our model to the current pandemic to show how government mandated quarantines might affect shopping behavior and returns to working capital investments. Section 8 concludes.

2 Data

Our analysis uses data from three main sources, the Nielsen Consumer Panel (NCP), the Nielsen Retail Scanner Panel (NRP), and the Survey of Consumer Finances (SCF).

2.1 Nielsen Consumer Panel (NCP)

The [Nielsen Company Consumer Panel \(2004–2014\)](#) consists of a long-run panel of nationally representative American households in 52 metropolitan areas. The goal of the NCP is to measure the detailed shopping behavior of American households while linking this data to household characteristics like household income, composition, age, and gender. Using bar-code scanners and hand-coded diary entries, participants are asked to report all spending on household goods that they engage in and also to detail information about the retail location that they visited in a given trip. Nielsen uses monetary prizes and continual engagement with panelists to try to maintain high levels of continued participation and limit attrition from the sample.⁷

The NCP is constructed to be a representative sample of the US population and fresh demographic information about participants is obtained each year. Nielsen maintains high quality data with regular reminders to participants that prompt them to report fully, and will remove non-compliant households from their panel. [Broda and Weinstein \(2010\)](#) provide a more detailed description of the NCP. [Einav, Leibtag and Nevo \(2010\)](#) perform a thorough analysis of the NCP, finding generally accurate coverage of household purchases though having some detectable errors in the imputed prices Nielsen uses for a subset of goods. Overall, they deem the NCP to be of comparable quality to many other commonly-used self-reported consumer datasets.

The NCP primarily covers trips to grocery, pharmacy, and mass merchandise stores but also spans a wider range of channels such as catalog and online purchases, liquor stores, delis, and video stores. The types of goods purchased span groceries and drug products, small electronics and appliances, small home furnishings and garden equipment, kitchenware, and some soft goods. Almost all of this spending is done in-store. In our sample years, under 5% of spending is done online or via catalog purchase in these categories.

In this paper, we utilize data from the 2013 and 2014 NCP. In each year, there are over 60,000 unique households with millions of individual shopping trips and tens of millions of individual

⁷Around 80% of NCP households are retained in the sample from one year to the next.

product purchases. Overall, the NCP tracks a sizable amount of a household's spending on material goods. On average, we observe over \$393 of spending per month for each household. This ability to measure household spending at the good level is key to our ability to understand inventory management as well as mechanisms by which households save on a given shopping trip.

2.2 Nielsen Retail Scanner Panel (NRP)

The [Nielsen Company Retail MSR Scanner Data \(2006–2014\)](#) 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 (e.g. 602 instead of 60208). This data covers almost 100 retail chains with over 40,000 unique stores in over 350 MSAs across the country.

In general, the data span a wide range of the largest retailers in the grocery, mass merchandiser, drugstore and pharmacy, and other miscellaneous retail sectors. Within the store, the data provide a comprehensive view of products sold, with more than 2 million unique UPCs across 1,100 product categories. During these years, the database picks up about half of total sales in grocery stores and pharmacies and about 30% of sales in other mass merchandisers. In total, these data comprise over 10 billion transactions per year worth nearly \$250 billion.

2.3 Survey of Consumer Finances (SCF)

The [Survey of Consumer Finances \(1983–2016\)](#) contains detailed information on U.S. households' income and assets. We define a stock market participation measure that is equal to one if a household owns stocks or stock mutual funds outside of retirement accounts. Income is gross household income over the calendar year preceding the survey. Financial assets include checking accounts, savings accounts, CDs, mutual funds, bonds, stocks, and money market funds.

2.4 Food Safety and Inspection Service Foodkeeper Data (FSIS)

The [Food Safety and Inspection Service FoodKeeper Data \(2020\)](#) contains information on recommended food and beverage storage times. We rely primarily on this information to infer depreciation estimates for each Nielsen product module.

3 Household Inventory Management

3.1 Computing Household Inventories

We now turn to computing consumer goods inventories at a household level across our sample. To compute household inventories using the NCP data, several assumptions are necessary. Although it is possible for us to track the flow of purchases for different items over time, the initial inventory is not observed and must be imputed. The rate of consumption is also not directly observed.

The first assumption we make is that a given product's consumption (and thus rate of inventory depletion) is constant throughout the year, and that total consumption equals total spending. Consistent with this approach, we must also aggregate individual products to a broader level to better understand consumption of certain product types. For instance, if a household switches cereal brands or types of apples purchased, they are not necessarily stocking up on all brands at once,

but keeping consumption of that product type constant using a substitute product. That is, if a household buys a different brand of the same product for variety and we use product-level inventory calculations, it will look like the household has bought that product on only one occasion and smoothed consumption over the entire year.

In practice, the household was consuming these products one after the other, so consumption was only continuous at a higher level of aggregation. For this reason, we group individual products at a Nielsen “product group” level when computing inventories.⁸

Next, we set initial inventory for each product group to the level that ensures that inventory for that product group is never negative during our observed sample window. Then we sum over the categories to get total inventory in both dollars and quantities.⁹ If households’ true inventories in each product group do not hit zero at some point during the year, our measure of inventories will be lower than the true level of inventories.

The validity of both the constant consumption assumption and the initial inventory assumption depend on the level of aggregation. If the product categories are too narrowly defined, the constant consumption assumption will be violated and inventory will be overstated. If the product categories are too broad the second assumption will be violated and inventory will be understated. Our choice of product group code is motivated by our personal assessment of the validity of the assumptions. However, we show below that even with a more conservative aggregation choice household inventories are still substantial.

Next, we derive a formula for annual average inventory. While the dataset only includes observations for days on which the household shops, the formula properly takes into account the time between trips.

The average inventory held over the period from $t = 0$ to T is $\frac{1}{T} \int_0^T I_t dt$, where I_t is the level of inventory at time t . Inventory at time t reflects the time zero level of inventory I_0 , purchases made on trips between time 0 and time t , and the rate of consumption, c , which we assume to be constant:

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

t_1, t_2, \dots, t_{n_t} are the dates of the n_t shopping trips occurring between time 0 and time t . S_{t_j} is the value of purchases made on the j th trip.¹⁰

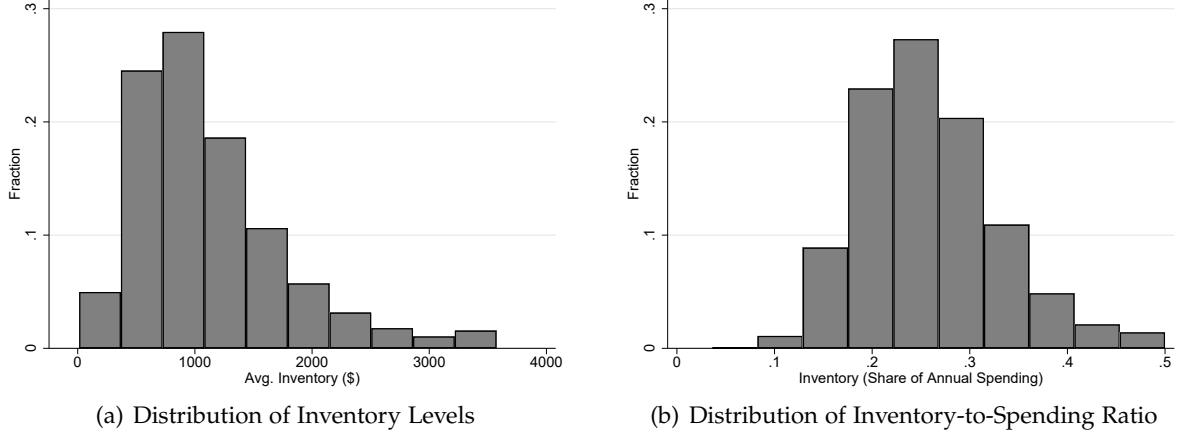
Next we compute the integral $\int_0^T I_t dt = I_0 T + \sum_{j=1}^{n_T} S_{t_j} (T - t_j) - c \frac{T^2}{2}$ and we divide by T to get

⁸This 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”. In general, if the assumed groups are too small, inventory will be overstated.

⁹We restrict attention to goods measured in “ounces.” This is the most common unit of measurement in the NCP and accounts for over half the UPCs. The other main unit of measurement is “count,” which does not allow for quantities to be compared reliably across different UPCs.

¹⁰In practice, some purchased goods deteriorate before they are consumed. Incorporating depreciation would mean, firstly, that spending reflects not only goods consumed, but also depreciation, and secondly, that the decline in inventories reflects both consumption and depreciation. Consequently, assuming that inventories decline in line with average spending is probably no less appropriate in the presence of depreciation.

Figure 1 – Observed Consumer Goods Inventory



Notes: We compute the average value of inventory for households in the NCP over 2013 and 2014. Panel (a) plots shows how the average inventory level varies across households. Panel (b) plots the distribution of inventory as a share of the household's annual spending on goods covered by Nielsen.

an expression for the average inventory:

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

When applying the formula to the data, we compute average annual inventory, so with t measured in years we have $T = 1$. Assuming annual spending is equal to annual consumption, annual average inventory is:

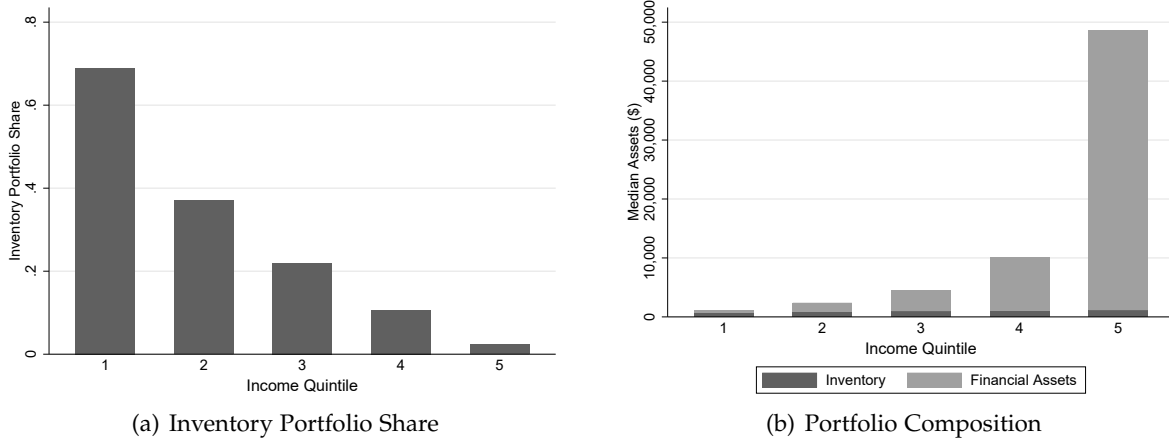
$$\text{Avg. Inventory} = \text{Initial Inventory} + \sum_{j=1}^N \text{Spending}_j \cdot \% \text{ of Year Left}_j - \frac{1}{2} \sum_{j=1}^N \text{Spending}_j. \quad (3)$$

N is the number of trips over the year, Spending_j is the value of products purchased on shopping trip j (i.e. trip size), and “% of Year Left $_j$ ” is the share of the calendar year remaining when trip j occurs.

With this approach (and over the product groups that the NCP data covers), the average amount of inventory for a household in the data is \$1,132. Figure 1 shows the distribution of our inventory measure across households. This measure of inventory naturally excludes inventory holdings in goods not covered by the NCP; most notably it excludes all large durable items like cars, furniture, most clothing and electronics. On the household balance sheet, such items would be classified as long-term physical assets – corresponding to “Property, Plant, and Equipment (PP&E)” on the corporate balance sheet – and are therefore not included in our definition of household working capital.

The average level of inventory is sensitive to the level of aggregation we assume. For transparency, we also report average inventory under alternative assumptions. The most conservative approach is to aggregate over all products the household consumes before backing out the initial inventory. This is likely to understate inventories, but still yields an average inventory value of \$511.

Figure 2 – Inventory Portfolio Share by Income



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 , Inventory_q . Finally, we compute the median level of financial assets held by the corresponding income quintile q in the SCF, $\text{Financial Assets}_q$, and also the corresponding inventory portfolio share, $\text{Inventory}_q / (\text{Financial Assets}_q + \text{Inventory}_q)$.

Nielsen includes a number of levels of product classification. Aggregating to the broadest product category, “department”, gives average inventory of \$726. Other possibilities include aggregating to “product module” or UPC. This yields average inventories of \$1,398 and \$1,870 respectively. In our opinion, these aggregation choices are likely to overstate inventories as the constant consumption assumption is probably inappropriate.

The result that households maintain a large stockpile of products on average is supported by Figure A.5, which shows that when households move to a new Zip Code they start to cut purchase quantities several months in advance. This is also accompanied by a drop in coupon usage (Figure A.6). The results continue to hold when restricting attention to staple items such as dried vegetables and grains, pasta and cereal.

Overall, inventory (even with durables excluded) is an important asset for many households. To show this, we compute household income quintiles using 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 2 shows the inventory portfolio share by income. For households in the bottom income quintile, inventories account for around 70% of assets. As income increases, inventory holdings grow more slowly than financial assets and the inventory portfolio share declines.

Table 1 shows that the inventory ratio is increasing in durability. It serves as a check on the magnitudes for our calculations of inventory levels. We manually assign each Nielsen product module a usable life in months, relying primarily on the FSIS data. Product life ranges from less

Table 1 – Relationship between Durability and Inventory Ratio


	(1)	(2)	(3)	(4)
Semi-durable	4.8*** (0.0)	5.3*** (0.0)	4.9*** (0.0)	5.2*** (0.0)
Durable	8.3*** (0.0)	8.5*** (0.0)	8.3*** (0.0)	8.3*** (0.0)
Potential Bulk Savings			8.9*** (0.1)	8.7*** (0.1)
Number of Trips (100s)			-3.7*** (0.0)	
Household FE		X		X
Number of Observations	1,781,712	1,781,712	1,652,696	1,652,696

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

$$\text{Inventory Ratio}_{i,c} = \beta_1 \text{Semidurable}_c + \beta_2 \text{Durable}_c + \beta_3 X_{i,c} + \text{Household FE} + \epsilon_{i,c} \quad (4)$$

Inventory Ratio is the ratio of household inventory to annual spending; see Figure 1(b). Columns 1 and 3 show results without household fixed effects. The base durability category is “Nondurable”. These are products which have a life of less than three months. We define semidurable items as those with a life of more than three months and less than one year. Coefficients are multiplied by 100. When computing inventory for this table, we assume that consumption is continuous within product group \times durability groups. Standard errors are clustered by household. * $p < .1$, ** $p < .05$, *** $p < .01$.

Figure 3 – Average Inventory by Department



Notes: This figure shows average inventory as a share of annual spending by product type using the NCP over 2013 and 2014. We compute the average inventory by household and department and then divide by the household’s average annual spending in that department. We drop observations where the household purchased less than one item per quarter on average in that department. Average product life is computed by assigning values to product modules based on the FSIS, and then aggregating to department weighting by product module expenditures.

than a week up to five years or more.¹¹ We consider products with a lifetime of less than three months to be non-durable. We define semi-durable products as those with a lifetime of at least three months and less than one year, and products with a lifetime longer than one year are considered durable. In Column 1, we see that households hold about an extra 2.5 weeks of spending in semi-durable products relative to non-durables, and an extra 4.5 weeks of spending in durable products relative to non-durables. Columns 2 through 4 show the relationship is robust to controlling for the number of shopping trips as well as household fixed effects.

Figure 3 shows average inventory levels by store department and their relationship to durability. The inventory ratio on the vertical axis is the ratio of inventory to annual household spending in that department. An inventory ratio of 0.1 corresponds to around 1.2 months of spending held as inventory, on average. Departments with high inventory are health and beauty, general merchandise, dry grocery, non-food grocery and frozen food. The departments with the low inventory are deli, packaged meat, fresh produce and alcohol. The horizontal axis shows log average product life. With the exception of alcohol, department inventory ratios are positively related to product life, consistent with Table 1. The relationship between inventory and durability is also non-linear. Conditional on products being storable, there is only a weak relationship between inventory ratios and product life.

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

3.2 Coupon, Deal and Bulk Savings

Households can obtain financial returns from investing in household working capital, by stocking up on goods that are on sale (i.e. “deals”), utilizing coupons, or by buying goods in bulk sizes at lower unit prices. This channel can act as a substitute to the channel identified by previous work that has focused on more frequent shopping trips to take advantage of lower prices (e.g. Aguiar and Hurst (2007)). In this way, people with a relatively high opportunity cost of time can obtain savings by allocating money to inventory instead of engaging in more frequent trips.


3.2.1 Coupon Savings

Coupon savings are relatively straightforward to measure in the NCP data as the total value of coupons used during a given shopping trip is directly reported. Average reported coupon savings are 6% of spending across households, though there exists substantial heterogeneity in coupon usage. Figure 4(a) shows how coupon savings increase with the ratio of inventory to total spending. That is, households that engage in more stocking up tend to utilize coupons more heavily than other households do. Consistent with Aguiar and Hurst (2007), and in line with the model we present in Section 5, use of coupons is also increasing the number of shopping trips as shown in Figure 4(b).

Table 2 shows estimates from the following regression specification, where i indexes households

¹¹Some products can be stored indefinitely (e.g. salt).

Figure 4 – Coupon Savings (%)



Notes: This figure is constructed using the NCP over 2013 and 2014. It shows bin scatter plots of the average coupon savings obtained for each decile of the average inventory ratio (Panel a, controlling for number of trips) and the average number of shopping trips per year (Panel b). For comparability with our measure of non-coupon deal savings and with the model, we compute coupon savings as a percentage of counterfactual household spending if the household had paid the average price reported for the same UPC and 3-digit zip code in the year the item was purchased.

Table 2 – Relationship between Inventory Ratio and Log % Savings

	(1) Coupon	(2) Deal	(3) Bulk
Log(Inv. Ratio)	120.3*** (5.6)	52.2*** (3.4)	-5.4*** (1.6)
Log(Potential Bulk)	-8.4*** (3.0)	13.5*** (2.8)	155.3*** (3.0)
Log(# Trips)	57.2*** (2.5)	18.4*** (0.8)	-4.6*** (0.8)
Income Group FE	X	X	X
Household Size FE	X	X	X
Number of Observations	97,533	111,903	118,491

Notes: This table reports estimates of regression specification (5) using data from the NCP over 2013 and 2014. The inventory ratio is the ratio of household inventory to annual spending; see Figure 1(b). The dependent variables in each column are log coupon savings, log deal savings, log bulk savings, and log total savings respectively. Potential bulk savings are the savings which would be obtained if the household bought the largest pack size available for each product. Savings are measured in dollars. Coefficients are multiplied by 100. Standard errors are clustered by household income group \times household size. * $p < .1$, ** $p < .05$, *** $p < .01$.

and Y_i are the different forms of savings from inventory management:

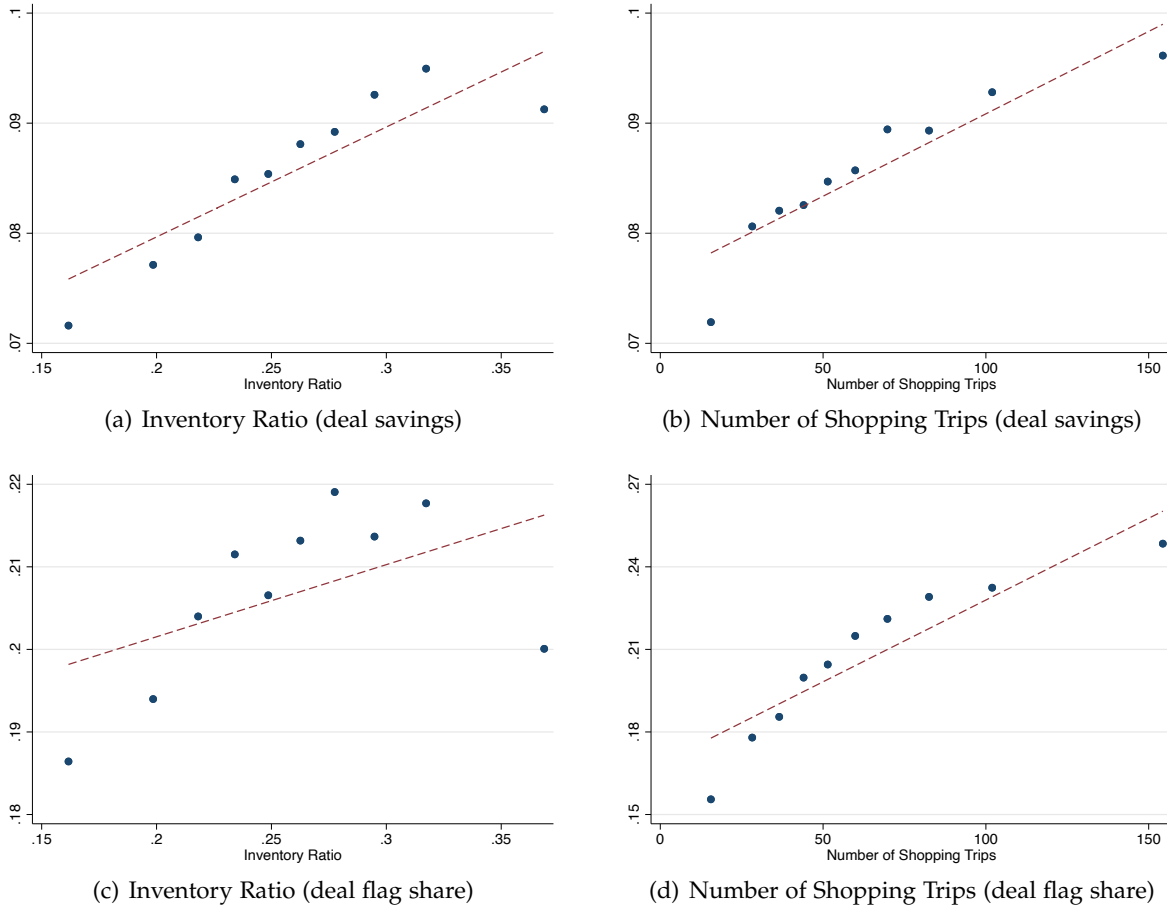
$$\log(Y_i) = \beta_1 \log(\text{Inventory Ratio}_i) + \beta_2 X_i + \text{Income FE} + \text{Household Size FE} + \epsilon_i. \quad (5)$$

Column 1 shows that increasing the inventory ratio by 10% is associated with a 12% increase in coupon savings.

3.2.2 Deal Savings

We construct savings from deals and general price reductions (or “sales”) by comparing the prices that households pay with the price they would pay if they engaged in “untargeted shopping” (or “inattentive shopping”) in their area. In particular, for each UPC-ZIP3 combination, we compute the average unit price for a year. This is an unweighted average from the retailer data, so it approximates the average price paid if the household randomly shopped across weeks and stores in the same ZIP3 (i.e. not targeting particular week-store combinations with a lower price).

Figure 5 – Deal Savings (%) and Deal Flag Share



Notes: This figure is constructed using the NCP over 2013 and 2014. The top row shows bin scatter plots of the average deal savings obtained for each decile of the average inventory ratio (Panel a, controlling for number of trips) and the average number of shopping trips per year (Panel b). The bottom row shows bin scatter plots of the average Nielsen deal flag share (excluding coupon events) for each decile of the average inventory ratio (Panel c, controlling for number of trips) and the average number of shopping trips per year (Panel d). We define deal savings as the % difference in price paid relative to the average price reported for the same UPC and 3-digit zip code in the year the item was purchased. We compute the average retail price using the Nielsen Retail Scanner Data. The Nielsen deal flag is equal to one for items the household considered to be on sale. While this includes coupons, to construct this chart we compute a measure of non-coupon deals, as we examine coupons separately in Section 3.2.1.

Using the consumer panel data, we compute the total amount actually spent by the household (excluding any coupon savings) and subtract the total amount they would have spent if paying the average per unit price. This number represents their dollar savings from sales. Average deal

savings are 11% of total observed household spending. This may seem low, but recall that savings are computed by comparing the price paid with the average price rather than the “full retail” price. Many goods are on sale relative to their full retail price and thus the average price a consumer would pay, even when shopping randomly for that item, is correspondingly lower than the full retail price.

In the top row of Figure 5, we show how deal savings increase with the ratio of inventory to total spending and also in the number of shopping trips. Column 2 of Table 2 shows that increasing the inventory ratio by 10% is associated with a 5.2% increase in deal savings.

In addition, the NCP data feature a “deal flag” equal to one if the household considered a purchased product to be a deal or on sale. The bottom row of Figure 5 uses the Nielsen deal flag, instead, finding consistent results for the inventory ratio, and a broadly flat relationship between self-reported deals and number of trips. While the deal flag includes coupon events, here we use a measure of non-coupon deals only as we consider coupons separately.

3.2.3 Bulk Savings

Finally, bulk savings represent savings obtained by buying a particular product in a large size or in a pack with multiple individual units at a lower per unit price. In general, each package size will have a separate UPC associated with it. Thus, to properly compute bulk savings, we are required to group UPCs associated with the same product. Unfortunately, this is not a straightforward exercise in the NCP data.

Our approach is to group products based on product module, brand, and common consumer name. Essentially we are trying to group otherwise identical products which are available in different sized packages.¹² We then compute pack size quintiles for each product that exists in multiple sizes. Looking across the quintiles of product size, we compute the average ZIP3 price associated with the second quintile.¹³ Bulk savings are then calculated as the per unit price actually purchased relative to the per unit price in the second quintile in the same ZIP3. On average, households save about 13% through buying in bulk.

There is considerable variation in potential bulk discounts across products. This means that some consumers may find it easier to take advantage of bulk discounts than others depending on the types of products and brands that they typically purchase. We construct a measure of potential bulk savings that we use to control for access to bulk discounts.¹⁴ This is defined as the difference between the average price in the highest quintile and the average price in the second quintile. For products where these quintile values are not defined due to limited dispersion, the potential bulk savings are set to zero.

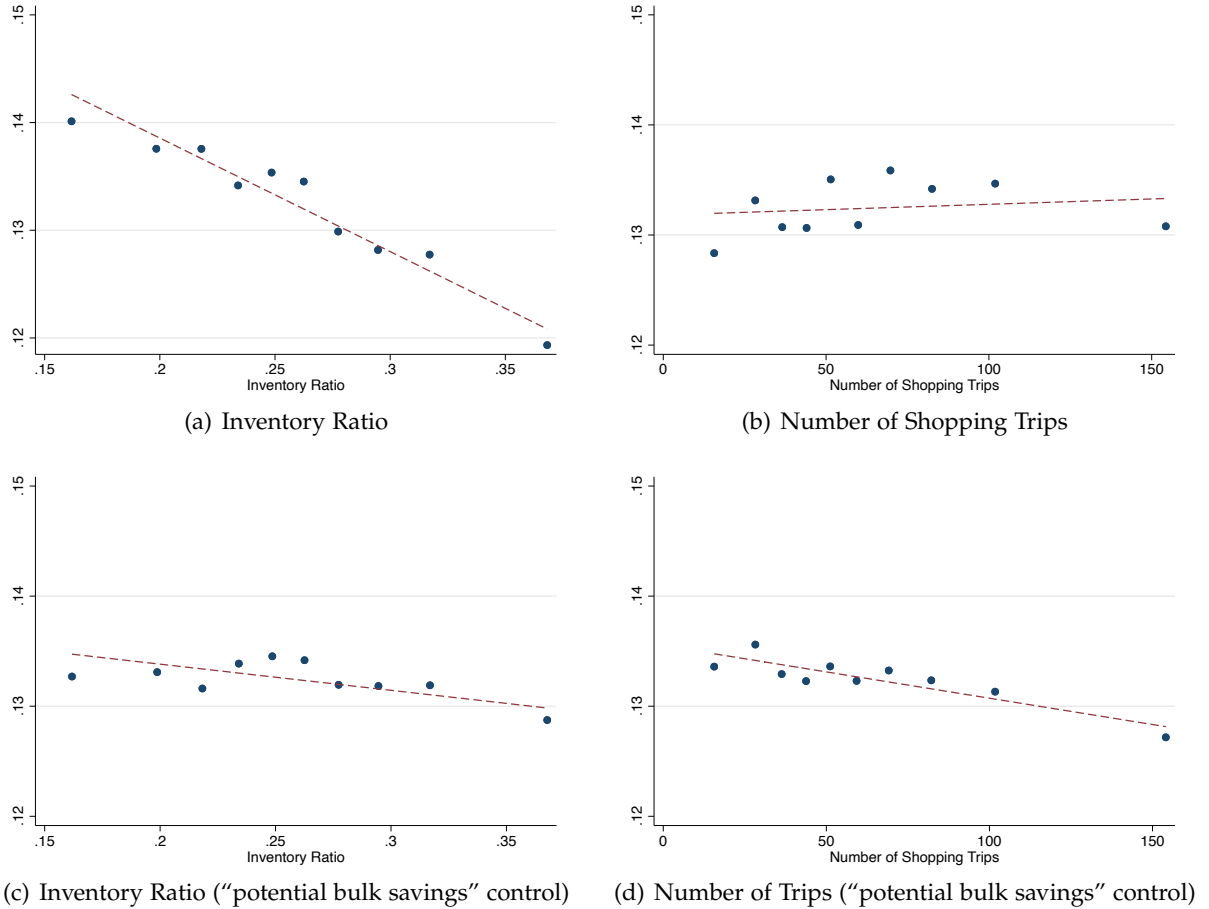
In general, we find that bulk savings at the household level are mainly driven by “potential bulk

¹²We manually inspect and drop combinations where this approach is problematic because the group is likely to contain products which are not identical. For example, we drop store brands because this group contains a large number of products that are likely to be different from each other. We also drop video products and nail polish – these modules contain a large number of products that are not easily substitutable because they are typically different colors or different films.

¹³We exclude the lowest quintile for most purposes because this typically represents a different use case (e.g. travel sizes of shampoo) rather than simply more units of a particular product.

¹⁴It is possible that a household’s product choice may also be driven by this bulk savings potential. Because of this possibility, we present results both with and without this control.

Figure 6 – Bulk Savings (%)



Notes: This figure is constructed using the NCP over 2013 and 2014. It shows bin scatter plots of the average savings from buying in bulk for each decile of the average inventory ratio (Panel a, controlling for number of trips) and the average number of shopping trips per year (Panel b). Savings from buying in bulk are computed by first grouping UPCs that are likely to correspond to the same product (for example, different pack sizes of Coca Cola). We then compute bulk savings as the difference in unit price between the quartile of pack size purchased and the second quartile of pack size for the same product. Panels c and d are conditional on "potential bulk savings." Potential bulk savings are the savings attainable if the household purchased the largest available pack size for each product. This is to control for the fact that holding their basket of goods or varieties fixed, some households have lower potential for bulk savings than others.

savings" (i.e. whether a household consumes products where sizable bulk discounts are available, as measured by the difference in average unit price between the top quintile pack size and the second quintile pack size). For many products, substantial bulk discounts are not possible. This phenomenon explains about 75% of the variation in bulk savings. Importantly, looking only at product size or quantity without also examining product type can lead to overestimating the potential bulk purchase savings. This is driven by the fact that cheaper products tend to come in larger pack sizes, on average. For instance, caviar and imported cheeses are expensive and typically come only in a small package size.

Figure 6 shows how bulk savings relate to the inventory ratio. We find that, unlike coupon and deal savings, bulk savings are basically unrelated to the inventory ratio. If anything, there is a negative relationship, but this largely disappears after controlling for potential bulk savings due to

Table 3 – Relationship between Log(Unit Price), Coupons, Deals, and Pack Size

	(1)	(2)	(3)	(4)
Coupon	-36.4*** (0.9)	-46.5*** (1.3)	-36.4*** (0.9)	-46.6*** (1.3)
Deal (Price < Local Avg.)	-26.6*** (0.9)		-26.7*** (0.9)	
Nielsen Deal Flag (Non-Coupon)		-8.1*** (0.7)		-8.2*** (0.7)
1 st Pack Size Quintile			12.6*** (2.0)	12.5*** (2.0)
3 rd Pack Size Quintile			-14.7*** (4.1)	-15.0*** (4.2)
4 th Pack Size Quintile			-22.5*** (2.9)	-22.8*** (3.0)
5 th Pack Size Quintile			-30.2*** (3.3)	-30.3*** (3.3)
Pack Size > 2 nd Quintile	-25.0*** (2.9)	-25.2*** (3.0)		
Product-Household FE	X	X	X	X
Number of Observations	44,809,003	44,809,003	44,809,003	44,809,003

Notes: This table is constructed using the NCP over 2013 and 2014. It reports estimates from the following regression specification, where i indexes transactions at the UPC level:

$$\log(P_i) = \beta_1 \text{Coupon}_i + \beta_2 \text{Deal}_i + \beta_3 \text{Bulk}_i + \text{Month FE} + \text{Product} \times \text{Household FE} + \epsilon_i. \quad (6)$$

P_i is the price paid by a household for a UPC on a particular trip. Coupon_i is an indicator equal to 1 if the purchase was made using a coupon and zero otherwise. Deal_i is an indicator equal to 1 if the transaction was classified as a deal and zero otherwise. Bulk_i is equal to 1 if the UPC purchased had a pack size greater than the second pack size quintile. Products are the product categories defined to measure bulk savings. Each product category contains multiple UPCs. Coefficients are multiplied by 100. Standard errors are clustered by product group code. * $p < .1$, ** $p < .05$, *** $p < .01$.

product choice. This suggests that households making bulk purchases do so because it is consistent with their “normal” purchases of that product. That is, they have a sufficiently high consumption flow because of large family size or just consuming large amounts of that product relative to other products. Their normal purchases may also be higher because they make infrequent trips, consistent with Figure 6(d). Column 3 of Table 2 shows that increasing the inventory ratio by 10% is associated with only a small 0.5% decrease in bulk savings.

Overall, we find that households do not periodically purchase a pack size much larger than what they ideally want in order to save money. This may be driven by the fact that a coupon and deal-based strategy offers at least as much if not better savings and also better align with desired quantities (e.g. depreciation can be an issue for many of these household goods if consumption is not sufficiently high). That is, it makes sense for big families to buy in bulk because it is at least somewhat cheaper and not costly in terms of depreciation, but for small households it may make more sense to use coupons, or buy items when they happen to be on sale. The approximately flat relationship between bulk savings and inventory is also consistent with our model results.

3.2.4 Product-level Analysis

In addition to analyzing the relationship between inventory and savings at the household level, we also run regressions at the product level to study how price and quantity purchased are related to whether the product was on sale or whether a large pack size was purchased.

Table 4 – Relationship between Log(Quantity), Coupons, Deals, and Pack Size

	(1)	(2)	(3)	(4)
Coupon	12.0*** (0.7)	18.7*** (1.1)	12.1*** (0.8)	19.1*** (1.1)
Deal (Price < Local Avg.)	3.3*** (0.5)		3.5*** (0.6)	
Nielsen Deal Flag (Non-Coupon)		15.2*** (1.9)		15.7*** (1.9)
1 st Pack Size Quintile			-24.1*** (4.7)	-24.2*** (4.6)
3 rd Pack Size Quintile			29.8*** (5.4)	29.7*** (5.3)
4 th Pack Size Quintile			58.6*** (10.1)	58.5*** (9.8)
5 th Pack Size Quintile			87.2*** (9.8)	87.7*** (9.7)
Pack Size > 2 nd Quintile	57.0*** (5.9)	57.0*** (5.7)		
Product-Household FE	X	X	X	X
Number of Observations	45,385,243	45,385,243	45,385,243	45,385,243

Notes: This table is constructed using the NCP over 2013 and 2014. It reports estimates from the following regression specification, where i indexes transactions at the UPC level:

$$\log(Q_i) = \beta_1 \text{Coupon}_i + \beta_2 \text{Deal}_i + \beta_3 \text{Bulk}_i + \text{Month FE} + \text{Product} \times \text{Household FE} + \epsilon_i. \quad (7)$$

Q_i is the quantity of a UPC (in OZ) purchased by a household on a particular trip. See Table 3 for more details. Coefficients are multiplied by 100. Standard errors are clustered by product group code. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 3 shows that each channel reduces the unit price of purchased goods. All specifications include product-by-household fixed effects. The average discount associated with a coupon event is about 36%. For deal events, it is about 27% (and approximately 8% when using the self-reported deal flag). For bulk savings, we use an indicator for purchases of products greater than the 2nd size quintile in columns 1 and 2, and size quintile indicators in columns 3 and 4. Savings when buying a larger pack size are about 25%.

In Table 4, we look at how quantity purchased is related to the different discount events. Using a coupon or getting a deal according to the Nielsen definition is associated with buying around 10-20% more on average. That is, households do indeed engage in “stocking up” when confronted with a deal or with lower prices. Pack size is very strongly related to buying more at the product level, almost by definition, though bulk savings are not positively related to total household inventory holdings, perhaps because these households tend to have higher consumption of such bulk-purchased products. Buying a pack size larger than the second quintile is associated with a

57% increase in quantity purchased.

4 Gross Savings from Inventory Management

By setting aside working capital, the household can reduce the average price it pays for consumer products. In order to understand the implications for other investment or borrowing behavior, we want to know the marginal financial return to allocating additional funds to working capital. We compute marginal returns using a calibrated model which we describe in Section 5. However, it is also possible to learn something about the potential returns to working capital directly from the data. In this section, we measure the relationship between average inventory and in-store savings. We refer to this as a gross return as it does not incorporate trip fixed costs or depreciation costs. The gross return is higher than the net return we compute using the model. We also use changes in observed inventory in the place of working capital, because we do not observe cash directly in the NCP, which is part of household working capital.¹⁵

In the previous section, we described how each of the different types of savings is computed and we illustrated the relationship between the inventory ratio and the percentage savings of each type. Now we show how the total savings households obtain on their purchases are related to the amount of inventory they hold. We can observe these amounts directly in the data using the aforementioned definitions, but the amount of inventory that a household holds varies for a range of reasons.

One of these is related to potentially endogenous trip frequency, where increasing inventory can decrease the number of shopping trips and thus the incurred trip fixed costs. For this reason, we hold the number of shopping trips constant in our calculations (e.g. a household might have large inventories of specific items based on what was on sale, but there are still other items they need to shop for). In general, the benefit of reducing fixed costs is only relevant at relatively low levels of inventory and, at the margin, what matters is how much the additional inventory can reduce the price per unit.¹⁶

Figure 7 illustrates the relationship after controlling for the number of shopping trips. In-store savings increase with inventory. On average an increase in the inventory ratio of 1 percentage point is associated with an increase in the savings ratio of 0.34 percentage points. This suggests a gross return of around 34 per cent over the range of inventory values we observe in the data. Figure A.4 shows that similar results are obtained when using Nielsen departments to compute average household inventories rather than product groups. For our final calculations of financial returns to household working capital investment, we will use the calibrated model in Section 5.


5 A Model of Optimal Household Inventory Management

Next, we use the NCP to calibrate a model of optimal household inventory management. We then use the model to compute the (net) marginal returns to household working capital investment,

¹⁵As we show in the model in Section 5, changes in household working capital are closely related to changes in average inventory. Conditional on holding a modest amount of household working capital, the model predicts that cash holdings display little relationship to the total amount allocated to household working capital. Instead, the additional household working capital is reflected in higher inventory holdings.

¹⁶In the context of the model presented in Section 5, it is possible for us to compute returns incorporating the trip fixed cost explicitly.

Figure 7 – In-store Savings and Inventory



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

taking into account holding costs due to depreciation and trip fixed costs, which are not directly observed in the data.

The model incorporates two types of savings: buying in larger quantities (“bulk”) and buying items on sale (“deals”). This essentially drives two key relationships between unit prices and shopping trip frequency. Buying in bulk relates directly to the size of the trip (i.e. the amount spent per trip) and buying items on sale relates directly to the frequency of the trip (i.e. more frequent trips yield on average more items on sale for a given trip size). Although the NCP distinguishes between deals and coupons, in the model we include both as “sale events” in the second category of savings.

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

Because buying large quantities reduces trip frequency and the ability to take advantage of sales, there is a trade-off between the two types of shopping policies to reduce the average unit price. In general, depending on various parameters (amount of household working capital, depreciation rate, shopping trip fixed cost, frequency and magnitude of sales, etc.), households may prefer one shopping policy over the other.

5.1 The Household's Problem

The household's problem is to minimize the cost of providing a monthly consumption flow of C , subject to an inventory constraint. For simplicity, we assume that the flow of consumption is constant both between trips and across trips.¹⁷ The cost per trip can be decomposed into two components – a fixed cost (e.g. the opportunity cost of time spent shopping) and a variable component which depends on the quantity of products purchased. The effective price per unit depends on the quantity purchased (bulk savings) and also on the household's choice of bargain-hunting policy (deal savings).

In the model, households consume goods with varying degrees of storability. Allowing for varying storability is important for matching the data and is consistent with what we observe in the NCP. Perishable goods are important for generating a realistic trip frequency, while storable goods allow us to simultaneously match the value of inventories. Appendix Figure A.3 shows that the majority of spending in the NCP is on products with a lifetime of either less than one month (23%) or more than one year (around 55%). Of the goods lasting more than a year, some can be stored for many years (e.g. some tinned products and cleaning products).

The household's problem is to choose the time between trips Δ (measured in months) and bargain-hunting policies $m_l \in \mathbb{N}_0$ for each good indexed by its level of perishability, l . The policy variable m_l is the maximum number of trips in advance that a household is willing to purchase and store an item with perishability level l . All goods are purchased according to the same trip schedule determined by Δ , but households are allowed to choose a distinct value of m_l for each type of good. This is described in more detail in Section 5.5.¹⁸

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

$$V(\bar{I}; \theta) = \min_{\Delta, \{m_l\}} \frac{k + \sum_l P_l(\Delta, m_l) S_l(\Delta)}{\Delta} \quad (8)$$

subject to:

$$\sum_l I_l(\Delta, m_l) \leq \bar{I}. \quad (9)$$

k is the shopping trip fixed cost. $P_l(\Delta, m_l)$ is the effective price per unit, taking into account bulk discounts, sales, and holding costs associated with setting $m_l > 0$ (i.e. stocking up in advance). $S_l(\Delta)$ is the quantity required immediately following a trip to satisfy consumption flow C until the next trip occurs.¹⁹ Section 5.3 describes how the trip interval Δ and the trip size S_l are linked given the requirement that inventory levels neither grow without bound, nor hit zero prior to the

¹⁷This assumption can be relaxed. For the CES case, see Baker et al. (forthcoming).

¹⁸We do not allow households to set different values of Δ for different goods. Although setting different values of Δ allows households to reduce depreciation costs, this is more than offset by the increase in trip fixed costs associated with maintaining multiple trip schedules, and so households prefer to buy all goods on the same trip. For a more detailed explanation of this tradeoff, see Bartmann and Beckmann (1992).

¹⁹When calibrating the model in Section 5.8, we normalize prices such that the effective unit price equals one on average for untargeted or inattentive shopping ($m_l = 0$) and when purchasing the "standard" pack size of each product ($S_l = \hat{S}_l$ and thus $\Delta = \hat{\Delta}$), i.e. $P(\hat{\Delta}, 0) = 1$, which is therefore the price of one physical unit of S_l .

next shopping trip. With this restriction, the household cannot choose Δ and S_l independently. Consequently, when the household chooses a trip interval, this directly implies a trip size.

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

The effective price function $P_l(\Delta, m_l)$ reflects bulk discounts and the net effect of bargain hunting (i.e. incorporating additional holding costs occurred due to stockpiling) and is characterized in Section 5.5.

At a given point in time, a portion of household working capital will be held as cash and the rest will be held as stored inventory goods. The inventory constraint means that the value of stored goods cannot, at any point in the shopping cycle, exceed the amount of assets set aside for managing inventory. The maximum inventory holdings occur immediately following a trip and at this point in time 100% of household working capital is held as stored inventory goods.²⁰

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

Ultimately, we are interested in the relationship between the dollar amount invested in household working capital and the dollar value of savings. In order for a particular shopping strategy to be feasible, the level of inventory immediately following a trip must not exceed the amount of household working capital \bar{I} . We will solve the problem at different levels of \bar{I} , and use this to compute the return to “investing” in household working capital (i.e. increasing \bar{I}). The investment payoff will be the reduction in the cost of providing the household’s exogenous consumption stream.

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

²⁰In this model, we abstract from other motives to hold cash, such as precautionary motives in response to income uncertainty or speculative motives to take advantage of time variation in investment opportunities (i.e. time-varying sales).

5.2 Implications for Portfolio Choice

In the model, working capital and consumption are exogenous. The model should be considered as one component of a higher-level problem in which the household chooses consumption and allocates assets to several investments, including working capital. Our model focuses on the choice of shopping strategy to minimize the cost of supplying a given consumption flow. Knowing how our model fits into this higher-level problem is important for understanding the implications for stock market participation. Therefore, before solving the shopping problem we demonstrate how our model fits into a static portfolio choice problem.

We consider the effect of working capital on the cost of supplying consumption to be analogous to interest earned on an investment. For simplicity of exposition, we assume the portfolio choice problem can be solved independently of the consumption problem.²¹ Assume the household has access to three investment opportunities: working capital, \bar{I} ; a risk-free bond, f ; and a risky asset m , 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:

$$\max_{\alpha_{\bar{I}}, \alpha_f, \alpha_m} = \int_{r_m} U\left([1 + \alpha_{\bar{I}}r_{\bar{I}}(\alpha_{\bar{I}}w) + \alpha_f r_f + \alpha_m r_m]w\right) dF(r_m) \quad (10)$$

subject to:

$$\alpha_{\bar{I}} + \alpha_f + \alpha_m = 1. \quad (11)$$

$\alpha_{\bar{I}}$ is the share of initial wealth w allocated to working capital, α_f is the share allocated to the risk-free bond with return r_f , and α_m is the share allocated to the risky asset.²² $F(r_m)$ is the probability distribution of returns for the risky asset, with $E[r_m] > r_f$.

$r_{\bar{I}}(\alpha_{\bar{I}}w)$ 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, we show that for working capital the return depends on both the investment amount and the level of consumption. Note that $\alpha_{\bar{I}}w$ in this problem corresponds to \bar{I} in the shopping model:

$$\bar{I} = \alpha_{\bar{I}}w. \quad (12)$$

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

Assuming consumers are risk averse, they choose $\alpha_{\bar{I}} = 1$ as long as $r'_{\bar{I}}(w) \geq E[r_m]$ because working capital has a higher expected return and lower risk over this range than the risky asset,

²¹The marginal return to working capital is independent of consumption assuming that both k and \hat{S}_l are scaled proportionally. This assumption is likely to be reasonable in cases where the opportunity cost of time is increasing in consumption, and where increases in consumption are reflected in purchases of higher quality products rather than purchasing larger quantities of the same products. In general, the marginal return to working capital does depend on the level of consumption, and the portfolio choice problem and the intertemporal consumption problem would need to be considered jointly.

²²The average inventory portfolio share in Figure 2 is closely related to $\alpha_{\bar{I}}$. $\alpha_{\bar{I}}$ additionally includes a cash component, but inventory accounts for the majority of working capital in our model for realistic parameter values.

and because investing in inventory also dominates the risk-free asset since $E[r_m] > r_f$. In Section 6, we show that our calibrated model delivers sufficiently high marginal returns that this is the case for sufficiently low levels of wealth. At higher levels of wealth the optimal allocation depends on the utility function, but as long as $r'_l(w) > r_f$ 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 and the household working capital constraint (9) no longer binds. Consequently, α_l gradually declines as wealth increases.

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

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

$$S_l(\Delta) = \int_0^\Delta e^{\delta_l t} C_l dt = \frac{C_l}{\delta_l} (e^{\delta_l \Delta} - 1). \quad (13)$$

If the trip size does not satisfy this condition, either inventory will grow without bound or the household will run out of a product before the next shopping trip. When discussing the bargain hunting policy below we consider the possibility that the household might want to buy an amount greater than $S(\Delta)$ if an item is on sale. We assume that the household chooses a multiple of $S(\Delta)$ and for convenience refer to these as packs (i.e. purchasing n packs means purchasing an amount $nS(\Delta)$).

5.4 The Bargain Hunting Policy, m_l

The bargain-hunting policy m_l is the maximum number of trips in advance that a household is willing to purchase and store an item with storability level l . Here we describe the household's strategy for taking advantage of random sales.

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

The effective price is thus $e^{\delta_l \Delta} p_d$ (the price paid in store is p_d , but the household incurs additional holding costs which lead to the price paid being multiplied by $e^{\delta_l \Delta} \geq 1$).²³ If the household decides not to buy the product now, the effective price is the expected price, $E[p] = x p_d + (1 - x) p_f$.²⁴

²³We assume that working capital is held in a zero interest checking account when it is not held as physical inventory, and therefore depreciation is the only holding cost. In practice, there may be other costs, including foregone interest. Relaxing this assumption has little effect on the optimal trip interval as perishable good depreciation costs are an order of magnitude larger than the risk-free return.

²⁴Note that some of the purchased product would also depreciate even further before it is consumed, but that additional depreciation cost applies in both cases and cancels out. Depreciation costs over the period when the pack is being consumed are captured in S .

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

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

Note that setting $m_l > 0$ does not lead to stochastic fluctuations in trip size or effective price when aggregating across a large number of products. However, it does lead to an increase in holding costs for a given trip interval Δ , because products will be bought in advance of when they are actually required for consumption. Increases in m_l will also increase inventory I , holding Δ fixed.²⁵

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

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

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

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

The expected price paid in store given the bargain-hunting policy m_l is therefore:

$$E[p|m_l] = p_f(1 - x)^{m_l+1} + p_d x \sum_{t=0}^{m_l} (1 - x)^t. \quad (14)$$

Note that the expected price with untargeted or inattentive shopping, $m_l = 0$, is $E[p|0] = xp_d + (1 - x)p_f = E[p]$. As the value of m_l increases, the probability that the item is purchased on sale approaches 1 and the expected price paid approaches p_d . Hence, ignoring shopping trip fixed costs, households in the model would optimally shop continuously and buy everything on sale.

²⁵If we allowed for additional complexity in the model, increases in m_l should increase the potential for bulk discounts at a given pack size because the pack size is spread over fewer products on any given trip. We abstract from this possibility for simplicity.

In practice, prices are clearly not independently distributed when shopping occurs at a very high frequency. Upon revisiting the store an hour later, prices are likely to be unchanged. However, at shopping frequencies observed among households in the Nielsen data (that do take into account shopping trip fixed costs; see [Baker et al. \(forthcoming\)](#)), and in our calibrated model, independence is a reasonable assumption.

5.5.2 Adding holding costs and bulk discounts

We factor the unit price P_l into two parts, the bulk discount as a function of the trip size, $b(S_l)$, and the shopping discount function arising from shopping strategy m_l . We begin with the latter.

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

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

Averaging over the $m_l + 1$ packs purchased gives $\frac{1}{m_l + 1} \sum_{i=0}^{m_l} e^{i\delta_l\Delta}$ per pack. In general, if the previous sale before this trip was $t + 1$ periods ago, the holding cost factor associated with the trip is $\frac{1}{t+1} \sum_{i=0}^t e^{(m_l-i)\delta_l\Delta}$. We compute the average effective price per unit using the probabilities from Section 5.5.1:

$$P_l(\Delta, m_l) = b(S_l) \cdot \left[p_f(1-x)^{m_l+1} + p_dx \sum_{t=0}^{m_l} (1-x)^t \frac{1}{t+1} \sum_{i=0}^t e^{(m_l-i)\delta_l\Delta} \right] \quad (15)$$

For simplicity, we assume that the bulk discount function $b(S_l)$ is applied directly the pack size $S_l(\Delta)$. Households can increase S_l , and take advantage of bulk discounts, by shopping less frequently (see Section 5.3). Bulk discounts therefore tend to raise the trip interval Δ .

We specify the bulk price discount function b to match bulk discounts observed in NRP data using the following functional form, implying that unit prices decline as the quantity purchased per trip S_l increases:²⁶

$$b(S_l) = \alpha + \beta e^{-\sigma \frac{S_l}{\hat{S}_l}}. \quad (16)$$

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

²⁶We compute the average unit price for each UPC over 2013 and 2014.

The function matches the data well in several respects: unit prices decay exponentially with pack size and converge to some level above zero. As pack sizes become very small, unit prices increase but do not become arbitrarily large. We normalize the price in the model to equal 1 when purchasing the standard pack size ($S_l = \hat{S}_l$) and in the absence of targeted deal shopping ($m_l = 0$). This means that α is interpreted as one minus the maximum % savings which can be obtained from buying in bulk.

To set \hat{S}_l , we solve the model without bulk discounts (i.e. with $b = 1$) and compute the optimal trip interval $\hat{\Delta}$. We then set $\hat{S}_l = S(\hat{\Delta})$. Because we calibrate the price distribution so that $E[p] = p_f x + p_d(1 - x) = 1$, the expected price per unit of S_l in the model is normalized to 1 for households purchasing the standard pack size and using an untargeted shopping strategy, i.e. $P_l(\hat{\Delta}, 0) = 1$. The calibration is described in detail in Section 5.8.

5.6 The Household Working Capital Constraint

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

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

$$I(\Delta, m_l) = P_l(\Delta, m_l) S_l(\Delta) + \mathbb{1}_{\{m_l > 0\}} S_l(\Delta) p_d \sum_{t=0}^{m_l-1} x(1 - x)^t \sum_{i=1}^{m_l-t} e^{(i+t)\delta_l \Delta}. \quad (17)$$

That is, the expenditure on the trip, $P_l S_l$, plus the value of inventory accumulated on previous trips to be consumed after the current trip. If $m_l = 0$, inventory is just the current trip value. In this case, inventory hits zero immediately prior to the next trip, and average inventory is $P_l S_l / 2$. The probability mass $\sum_{t=0}^{m_l-1} x(1 - x)^t$ is equal to $1 - (1 - x)^{m_l}$, where $(1 - x)^{m_l}$ is the share of goods for which no inventory remains at trip time.

5.7 Solution Method

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

1. Define a grid over trip intervals Δ and bargain-hunting strategies $\{m_l\}$, where $m_l \in \mathbb{N}_0$.
2. For each possible combination of $(\Delta, \{m_l\})$, compute $S_l(\Delta)$, $I_l(\Delta, m_l)$, and the value of the cost function, $\frac{k + \sum_l P_l(\Delta, m_l) S_l(\Delta)}{\Delta}$.

Table 5 – Model Calibration

Name	Parameter	Calibrated value	Source/target
Trip fixed cost	k	\$4.85	Baker et al. (forthcoming) .
Deal probability	x	0.24	x is calibrated to match NCP deal share.
Full list price	p_f	1.064	p_f and p_d are jointly calibrated to match average discount size in the NCP and $E[p] = xp_d + (1 - x)p_f = 1$.
Deal price	p_d	0.794	
1 - max bulk savings %	α	0.85	α , β and σ are jointly calibrated to match the relationship between pack size and unit price in the NCP.
Bulk savings parameter	β	0.89	
Bulk savings parameter	σ	1.80	
Monthly consumption	C	464.83	C and δ_0 are jointly calibrated to match average annual spending and inventory in the NCP.
Storable depreciation	δ_0	0.0026	
Perishable depreciation	δ_1	2.88	Sets expected expiry to average expiration date for products lasting < 1 month (0.347 months)
Storable good share	s_0	0.769	NCP expenditure share of goods with expiration date ≥ 1 month

- Find the values of $(\Delta, \{m_l\})$ which minimize the cost function subject to the household working capital constraint.

5.8 Calibration

We calibrate the model by choosing θ to match a number of data moments summarized in Table 5. We set the fixed cost per shopping trip to $k = \$4.85$ as in [Baker et al. \(forthcoming\)](#). As discussed there, this is consistent with a hourly before-tax reservation wage of \$7.12 to \$10.69.

To calibrate p_f and p_d , we estimate the average price drop associated with a discount event in the NCP.²⁷ Given an estimated log price difference of 0.293, we set $\frac{p_f}{p_d} = e^{0.293} = 1.34$. We pin down p_f and p_d by normalizing the expected price achieved using an untargeted shopping strategy without bulk discount $E[p] = xp_d + (1 - x)p_f = 1$.

To calibrate the bulk discount function $b(S_l)$, we first choose the value \hat{S}_l which corresponds to the standard pack size for which there is no bulk price discount, i.e. $b(\hat{S}_l) = 1$. We set $\hat{S}_l = S_l(\hat{\Delta}) = \frac{C_l}{\delta_l}(e^{\delta_l \hat{\Delta}} - 1)$, where $\hat{\Delta}$ is the optimal trip interval in the model without bulk discounts.²⁸ Because we also set $E[p] = 1$, this means that the expected effective price per unit of S_l in the model

²⁷We define a discount indicator D_u which is equal to 1 when a product is purchased either with a coupon, or at a price which is lower than the annual UPC-ZIP3 average price. We estimate the average log discount using $\log \text{Price}_u = \beta_1 D_u + \beta_2 \text{Bulk}_u + \text{Month FE} + \text{Product} \times \text{Household FE} + \epsilon_u$, where u indexes transactions at the UPC level.

²⁸The definition of a “standard” pack size is clearly somewhat arbitrary. Our assumption is that the second quartile of package sizes in the data correspond to the amount households in the model would want to buy if they were not influenced by the presence of bulk discounts. This assumption yields bulk savings which are similar to what we observe in the data.

is normalized to 1 for households purchasing the standard pack size and using an untargeted shopping strategy, i.e. $P(\hat{\Delta}, 0) = E[p] = 1$.

We calibrate the parameters of the function $b(S_l)$ by estimating relationship (16) with weighted least squares as follows. First, we prepare the Nielsen data by creating a new product ID as described in Section 3.2.3.²⁹ 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. We also want to express both prices and pack size relative to the standard pack size for that product, S_l/\hat{S}_l . To do this, we compute the average number of units in the second quartile of package sizes for each product, as well as the average price per unit.³⁰ We then divide by these second quartile averages.

In the model, we assume that households purchase the same multiple of the standard pack size across all the products with the same perishability. When households in the model take advantage of bulk discounts by increasing trip size, this necessarily coincides with reduced trip frequency (holding consumption fixed).³¹

In the data, a large proportion of spending is accounted for by products which have only limited bulk savings potential (for example, it may not be possible to purchase a pack size more than 1.5 times the standard pack size, leading to households paying the same unit price even though the amount purchased has increased). This means potential bulk savings would be overstated if we were to estimate the relationship between pack size and price without further adjustments, as the relationship at higher pack sizes would be based only on products for which extreme pack sizes are available.

We therefore aggregate the data to pack size group \times product, and then make sure the dataset is balanced (i.e. every product has a non-missing price for each pack size group). For products where large pack sizes are not available, we assume the household obtains the unit price associated with the largest available pack size. We compute total expenditure for each product and use this to weight our regressions. We estimate the following relationship for different values of σ ,

$$\text{Price} = a_0 + a_1 e^{-\sigma \text{Units}}, \quad (18)$$

and choose σ to maximize the within R-squared. This yields $\hat{\sigma} = 1.80$. The estimates of a_0 and a_1 from the same specification are $\hat{a}_0 = 0.84$ and $\hat{a}_1 = 0.88$ respectively. The results are shown in Appendix Table B.1. For the model, we normalize the price of the standard pack size to one, and the price of other pack sizes reflect percentage deviations from the standard pack size. We therefore calibrate α and β using $\alpha = \frac{\hat{a}_0}{\hat{a}_0 + \hat{a}_1 e^{-\hat{\sigma}}}$ and $\beta = \frac{\hat{a}_1}{\hat{a}_0 + \hat{a}_1 e^{-\hat{\sigma}}}$. Figure 8 compares the bulk discount function we use in the model with the corresponding relationship in the data.


We calibrate C and δ_0 jointly to match the average monthly spending and inventory in the NCP. The average monthly spending on goods covered by Nielsen is \$393, the average level of household

²⁹As explained in Section 3.2.3, we restrict attention to products measured in ounces, which is also the most common unit of measurement. The number of ounces associated with a given UPC is computed as `multi × size1_amount`.

³⁰We compute the average per unit price as the expenditure weighted annual average price for that UPC in the NRP.

³¹If households did not scale their purchases of all products in the same way, this raises the possibility that they run out of some items before the next trip. For simplicity, we choose to abstract from this.

Figure 8 – Bulk Calibration



Notes: The solid line shows the average per unit price relative to the second pack size quintile (for the same product), weighted by spending, using the NCP over 2013 and 2014. The dashed line shows the relative price we assume in the model: $\text{Price} = \alpha + \beta e^{-\sigma \text{Units}}$, where $\alpha = 0.85$, $\beta = 0.89$, and $\sigma = 1.80$.

inventory is \$1,133. Because it is not in general possible to match both quantities exactly, we use values of the parameters that minimize the sum of squared percentage deviations between the model and the data. The corresponding model values are monthly spending of \$397 and average inventory of \$1,123.

As discussed above, the quantity units in the model are normalized so that the price per unit in the absence of any deals or bulk discounts is \$1. The calibrated value of C is 465 units. The calibrated value of x is 24%. This means the discount price is observed about once every 4 trips. To calibrate x , we use the deal flag from the NCP, which is equal to one for purchases where a coupon was used, or where the household considered the item to be on sale. We estimate the following relationship using the NCP:


$$\text{Deal \& Coupon Share}_{h,t} = \alpha + \beta_1 \text{Inventory Ratio}_{h,t} + \beta_2 \text{Trips}_{h,t} + \epsilon_{h,t} \quad (19)$$

We then set x equal to the predicted deal and coupon share for a household with an inventory ratio equal to $0.5 \times \frac{1}{\text{Trips}_{h,t}}$, where $\text{Trips}_{h,t}$ is the number of trips made by household h in year t . This is the average inventory ratio a household would have if they did not engage in strategic deal shopping.

We calibrate perishable depreciation δ_1 and the storable share s_0 by combining USDA information on product storage with information on households' purchases from the NCP. We define perishable goods ($l = 1$) as products which have expiration dates of less than one month. Using the USDA information, we manually assign a time to expiry for each Nielsen product module.

We assume exponential product depreciation and calibrate δ such that the average expiry time, $\frac{1}{\delta}$, matches the expiry time given by the USDA for the relevant set of products. The average time to

Figure 9 – Model Fit: Deal and Bulk Savings (%)



Notes: We evaluate bulk and deal savings in the model for different values of \bar{I} . The variation in average inventory levels on the x-axis is therefore generated by a relaxation of the working capital constraint.

expiry for perishable goods is 0.347 months, giving $\delta_1 = 2.88$. Given an average expenditure share of storable goods in the NCP of 76.9%, we set $s_1 = 0.769$.³²

5.9 Model Fit

We assess the fit of the model by looking at statistics we do not directly target in the calibration: the relationship between the inventory ratio and % deal savings, the relationship between the inventory ratio and % bulk savings, and the number of days between trips.

We compute the percentage savings obtained by buying in bulk, or taking advantage of deals. Comparing Figure 9(a) with Figures 6(c), we can see that the model generates a similar level of bulk savings to the data and also a weak relationship with the inventory ratio. Figure 9(b) shows the relationship between deal savings and the inventory ratio in the model. Consistent with the data, there is a stronger relationship between deal savings and the inventory ratio (see Figures 4(a) and 5(a)). The overall level of deal savings is also comparable. Combining deal and coupon savings in the data leads to total savings of just over 15% at high inventory ratios, which is a little higher than the model but broadly comparable.

The optimal trip interval in the model of just over one week is also consistent with the data. The median time between trips to the same grocery retailer in the NCP is 7 days.

6 Financial Returns to Household Inventory Investment

Solving the optimization problem (8) yields the average monthly cost $V(\bar{I})$ of supplying consumption flow C . To compute the return to household working capital, we compute this cost at each level of household working capital \bar{I} .

In principle, we can then compute the marginal return as $V'(\bar{I})$, providing a net return measure which incorporates not just the price paid in store (as in Section 4) but also shopping trip fixed

³²One way to think about this assumption is that the expiry time for a product has a Poisson distribution with parameter δ . While this means there is some chance that items will last a lot longer than the average expiry time, it does still capture the high cost of storing perishable items and generates a trip frequency consistent with the data.

costs and depreciation costs. In practice, because m_l is discrete, the cost function is not smooth. 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. In the tables below we use an increment of 5 per cent of annual spending. We multiply by 12 to convert monthly to annual returns.³³

Table 6 – Financial Returns to Household Inventory Investment

Working Capital Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings: Deal Bulk		Interval Δ^* (Months)	m_0^*	Net Return (%)
0.05	133.52	104.30	9.46	10.94	0.25	4	72.30
0.10	370.96	101.54	12.66	11.05	0.25	8	21.85
0.15	604.66	106.25	13.48	11.69	0.27	11	6.89
0.20	839.45	107.53	13.86	11.86	0.27	14	1.42
0.25	1069.60	107.46	14.03	11.86	0.27	17	0.00
0.30	1069.60	107.46	14.03	11.86	0.27	17	0.00

Notes: Working capital ratio is the ratio of household working capital in \$ to average annual spending in the data ($\$393 \times 12 = \$4,716$). Maximum cash is the cash held immediately prior to a shopping trip. This is also the same as the total price paid in store in \$, $\sum_l P_l \cdot S_l$. Immediately following a trip, cash in the model is equal to zero (because there are no other motives to hold cash, such as precautionary or speculative motives). Average cash holding in the model for purchases of goods covered by Nielsen is therefore approximately half of the maximum holding, and the maximum cash holding is equal spending per trip. Minimum inventory is the inventory level immediately prior to a shopping trip. The maximum level of inventory is equal to the sum of minimum inventory and maximum cash and is reached immediately after a shopping trip. Deal savings are % savings of annual spending due to buying an item on sale. Bulk savings are % savings of annual spending due to buying a larger pack size. Interval Δ^* is the optimal length of time between trips measured in months. m_0^* is the optimal deal shopping strategy for the most storable goods, subject to the household working capital constraint. The net return incorporates not only in-store savings but also depreciation and trip fixed costs.

Table 6 shows how increasing the maximum household working capital \bar{I} affects inventory and cash holdings, as well as the different sources of savings households are able to achieve. When the amount of funds allocated to household working capital is low, the household is restricted in its ability to take advantage of deals. This is because stockpiling products well in advance of when they are needed (i.e. a large m_l) is working capital intensive. Low levels of household working capital investment therefore constrain households to choose a low value for m_l (conditional on the fixed trip cost being non-trivial).

As the household working capital investment is increased, households choose progressively higher values of m_0 .³⁴ Under a deal-focused strategy households tend to shop more frequently, which reduces their trip size and tends to weigh on their bulk savings. With a deal-focused shopping strategy households make small and frequent trips on which they buy only a subset of their consumption bundle. They purchase only goods that are on sale, or goods for which inventory has been run down to zero.

³³We assume that the working capital investment remains fixed at \bar{I} throughout the year. Given that the marginal return is diminishing in \bar{I} , it is not appropriate to assume the proceeds can be reinvested at the same rate of return. To the extent that monthly returns are invested elsewhere and earn a positive return, the annual return will be larger than what we assume.

³⁴Because of the high depreciation costs associated with stockpiling perishable items, m_1^* is always equal to zero for the parameter values we consider here.

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 fixed. This force works in the opposite direction, pushing the trip interval up and raising bulk savings. Given the parameter values we use here, increasing the trip interval does not exert much downward pressure on m_0^* and deal savings. This is because the storable good has very low holding costs, and the perishable good is not worth stockpiling at the levels of \bar{I} we consider. Consequently we see the optimal trip interval increase slightly with working capital.

In general, model outcomes such as the marginal return, minimum inventory, trip interval and percentage savings need not be monotonic in \bar{I} . As discussed above, the cost function is non-smooth because m_l is discrete. Furthermore, relaxing the constraint may have either a positive or negative effect 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 simultaneously makes it more costly to buy items several trips in advance and can therefore lead to a reduction in minimum inventory. Alternatively, if the household uses the additional funds to buy items on sale well in advance, this tends to put downward pressure on trip size due to depreciation costs and reduces bulk savings.


At low levels of household working capital investment, the marginal return to additional investment is very high. We compute marginal returns using a working capital increment of 5% of annual spending. When household working capital is equal to 5% of annual spending, the marginal return is 72%. The marginal return gradually diminishes and reaches zero when household working capital is around three times monthly spending.

We obtain qualitatively similar results when using a more conservative value of inventory to calibrate the model. Table B.2 shows the results using an average inventory level of \$721. As discussed in Section 3.1, this is the average value of inventories we obtain when aggregating to Nielsen departments rather than product groups when computing inventories.

Figure 10 shows that the model's prediction is consistent with the data of more educated households. It shows that direct stock market participation (direct holdings of stocks and stock mutual funds) is closely related to the value of net financial assets. These educated households come closest to the rational, optimizing agents in the model. They likely also understand that they can obtain a substantial return premium from investing in risky financial assets and how to best achieve such expected excess returns. For these households, we see that stock market participation rates are indeed nearly zero at low levels of financial assets, where the returns from inventory management dominate stock returns. Once assets exceed around one month's income and the marginal return from inventory management decreases, participation in the stock market increases sharply.

The level of cash holdings predicted by the model should of course not match the level of cash holdings observed in the SCF 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 probably also applies to other goods not covered by the Nielsen data which also require additional cash holdings.

Figure 10 – Stock Market Participation by Education



Notes: This figure is constructed using the SCF over 2010, 2013 and 2016. Financial assets include checking accounts, savings accounts, CDs, mutual funds, bonds, stocks and money market funds. Measures of participation are based on directly held stocks or stock mutual funds, and do not include retirement accounts. We use the education of the spouse with the highest level of attainment. Each point corresponds to a quartile. The figure excludes households with a ratio of financial assets exceeding 50 per cent of annual income.

Households with less education tend to have a low level of direct stock market participation even at a high level of financial assets. As discussed in the introduction, such behavior can be rational in light of previous studies that have shown that households with low education tend to earn low returns on their stock market investments conditional on participation. For example, these low returns may be driven by under-diversification, local bias, excessive trading, or paying high management fees. Investing in household inventory offers these household an alternative investment opportunity with reasonably high returns at no or only low risk. Moreover, households can implement this investment themselves, without any need to delegate to a financial advisor, because they have a lot of experience with shopping.

7 Application: Household Inventory Management During a Pandemic

During the COVID-19 outbreak in 2020, household spending and shopping habits were impacted dramatically. COVID-19 was first seen in late 2019 in Wuhan, China, before spreading worldwide over the next several months. In the United States, the first case was identified on January 21, 2020, but was soon followed by increasing case loads across all regions of the country. With cases and death tolls mounting, various state and federal officials declared states of emergency and began to restrict international as well as domestic travel.

In particular, a number of states began to announce variants of “shelter-in-place” policies alongside restrictions on “non-essential” business activities. Many retail stores were to be closed to foot traffic and most restaurants were to only serve takeout or delivery food rather than allowing in-store

Table 7 – Effect of Shopping Trip Fixed Cost

Fixed Cost (\$)	Min. Inv. (\$)	Max. Cash (\$)	% Savings: Deal Bulk		Interval Δ^* (Months)	m_0^*
1.00	1006.25	100.97	14.15	11.15	0.26	17
2.00	1022.08	102.57	14.12	11.34	0.26	17
4.00	1053.76	105.82	14.06	11.69	0.27	17
8.00	1125.08	113.32	13.93	12.37	0.28	17
16.00	1154.43	126.45	13.67	13.17	0.31	16
32.00	1346.22	151.11	13.28	13.97	0.37	16
64.00	1521.79	195.07	12.57	14.48	0.45	15
128.00	1733.60	263.81	11.57	14.67	0.56	14

Notes: The fixed costs in the first column are the cost of making a trip to the store in \$. To isolate the effect of different fixed costs on inventory management, we supply the household with enough working capital such that constraint (9) is not binding and hence the implied net marginal return is zero in all rows. The remaining columns are described in Table 6.

dining. Even in locations where such policies were not put in place by local governments, many households began to adjust their daily routines to minimize the possibility of contracting COVID-19.

In the context of our model, we interpret the circumstances of and policy responses to this pandemic as an unexpected increase in the fixed cost of shopping. Using our model, we can illustrate the impact that this change would have on the shopping habits of American households.³⁵

We perform two counterfactuals with our model. First, we look at the same outcomes for households as trip fixed costs increase. Second, we investigate the impact of a high trip fixed cost of \$100 on the shopping habits of households who have varying levels of household working capital. We display our results of these experiments in tables taking the same form as our main model results in Table 6.

In Table 7, we allow fixed costs to vary from \$1 per trip to \$128 per trip (doubling with each increment). For comparison, our baseline value for a trip fixed cost was \$4.85. As fixed costs increase, households tend to increase the trip interval – a fixed cost of \$2–\$4 yields a household shopping around once per week while a trip cost of \$128 means that the household only visits the store around twice per month. The fact that the non-storable good share in our model is fixed at normal levels restricts the extent to which households can increase the trip interval. In practice, allowing for substitution away from these perishable products when trip fixed costs rise would allow for even larger reductions in trip frequency.

As trip intervals increase, the size of the shopping trip must increase to maintain levels of consumption over the longer intra-trip period. We find that the amount spent per trip (i.e. maximum cash holdings in Column 3 of Table 7) more than doubles in our range of fixed costs. Since it becomes optimal for households to increase purchase quantities, bulk savings also increase. In contrast, at low trip frequencies substantially more working capital is required to stock up on storable goods. Consequently, deal savings decline by about 20% in the range of fixed costs we analyze.³⁶

³⁵While online purchasing and delivery services saw increased usage during this period, the majority of dollars spent at retailers were done in-person. Our model examines only the in-person shopping response of these changes in trip fixed costs.

³⁶If the storable good were somewhat less storable than we assume, it is possible that deal savings may fall to zero as

Table 8 – Model Results with High Shopping Trip Fixed Cost ($k = \$100$)

Working Capital Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings: Deal Bulk		Interval Δ^* (Months)	m_0^*	Net Return (%)
0.05	29.50	208.44	3.05	14.44	0.44	1	163.72
0.10	254.81	220.96	8.28	14.56	0.48	4	48.97
0.15	483.32	228.17	9.93	14.60	0.50	6	22.59
0.20	726.95	224.64	10.97	14.59	0.49	8	10.05
0.25	894.69	238.01	11.13	14.63	0.52	9	8.29
0.30	1174.26	237.17	11.57	14.63	0.52	11	2.91

Notes: See the description in Table 6.

The relationship between fixed costs and inventory is in general non-monotonic and reflects two competing forces. First, with low fixed costs, households tend to make many small trips, buy only a subset of goods each trip and maintain a sizeable baseline level of inventory. As fixed costs increase, trip frequency declines and depreciation costs rise. This tends to reduce m_0^* . All else fixed, this would lead to a reduction in inventory; however, the lower trip frequency works to offset this. When trips are less frequent, a higher level of inventory is required to support a given value of m_0^* . Because the storable good has low holding costs, the reduction in m_0^* is modest in this case.

Figure 11 illustrates the relationship between number of trips, inventory ratio, and savings when the variation in number of trips is generated by variation in the fixed cost k . Figure 11(a) illustrates the non-monotonic relationship between the inventory ratio and the number of trips discussed above. The upward jumps in Figure 11(a) occur at points where m_0^* increases.


Figure 11 highlights why it is important to control for the number of trips a household makes when analyzing the relationship between inventory and savings using the NCP. In our model, variation in the inventory ratio generated by fixed costs is negatively related to deal savings (whereas variation generated by working capital is positively related to deal savings). Controlling for household trip frequency allows us to control for differences in the fixed cost of shopping across households and focus on the working capital channel.

Baker, Farrokhnia, Meyer, Pagel and Yannelis (2020) found some evidence for this sort of stockpiling behavior that might indicate higher levels of fixed costs. They use transaction-level financial data to highlight changes in household behavior as the COVID-19 pandemic began to spread in the United States. They show that in mid-March 2020, as shelter-in-place policies began to be announced, household spending rose substantially as they stockpiled retail and grocery goods. In locations with higher numbers of COVID-19 cases and earlier shelter-in-place policies (making retail trips more costly), spending increased by more than in other locations. In subsequent weeks, spending declined to record low levels as households ran down purchases and were unable to shop at many of their typical retail outlets.

We examine further variation in this setting in Table 8. Here we test the behavior of households who hold different levels of household working capital given elevated trip fixed costs of \$100. Given the large expenditure share of highly storable goods in the NCP, high trip fixed costs increase the

fewer trips make depreciation costs prohibitive. In our case this does not happen, because even with a trip interval as long as a month depreciation costs are still sufficiently low to justify stocking up when goods are on sale.

Figure 11 – Relationship between Inventory Ratio, Deal Savings and Number of Trips



Notes: We evaluate the inventory ratio and % savings in the model for different values of k . Working capital is set sufficiently high that the constraint does not bind. Number of trips is computed as $\frac{12}{\Delta}$. Variation in the number of trips is generated by variation in trip fixed costs.

returns to working capital. With a longer interval between trips, households need a substantial amount of working capital to cover the high costs associated with large trips, as well as stock up on storable goods in response to deals. Table 8 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.³⁷

Again, [Baker et al. \(2020\)](#) find some evidence to support this result. They find that households possessing higher levels of income and liquidity in the weeks leading up to shelter-in-place orders tended to stock up to a greater extent than households with lower levels of financial assets.

Overall, we find that our model predicts significant changes to household shopping behavior when trip fixed costs increase substantially. Moreover, “out-of-sample” evidence from the onset of the COVID-19 pandemic in the United States is consistent with our model’s prediction of increased trip costs during this period. Therefore, the level of household working capital is a significant driver of households’ shopping behavior.

8 Conclusion

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

³⁷This result is dependent on households consuming a highly storable good. In the absence of such goods, high fixed costs lead households to set $m = 0$ regardless of the level of working capital.

high returns from inventory management through several channels at low levels of inventory, but these returns decline rapidly as inventory levels increase. At low levels of inventory, the marginal return to investment in inventory strongly dominates stock market returns and it can even dominate some forms of borrowing costs such as credit card interest rates and fees.

Finally, we apply our findings to study how government regulations and quarantines can affect households' strategic shopping behavior and inventory management.

References

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino**, "Perception of House Price Risk and Homeownership," *NBER Working Paper No. 25090*, 2020.
- Aguilar, Mark and Erik Hurst**, "Life-Cycle Prices and Production," *American Economic Review*, 2007, 97 (5), 1533–1559.
- Anagol, Santosh, Shawn Cole, and Shayak Sarkar**, "Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market," *Review of Economics and Statistics*, 2017, 99 (1), 1–15.
- 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.
- Athreya, Kartik, Felicia Ionescu, and Urvi Neelakantan**, "Stock Market Participation: The Role of Human Capital," *University of Pennsylvania Working Paper*, 2015.
- Badarinza, Cristian, John Y. Campbell, and Tarun Ramadorai**, "International Comparative Household Finance," *Annual Review of Economics*, 2016, 8 (1), 111–144.
- Baker, Scott R., Robert A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis**, "How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic," *NBER Working Paper No. 26949*, 2020.
- , **Stephanie Johnson, and Lorenz Kueng**, "Shopping for Lower Sales Tax Rates," *American Economic Journal: Macroeconomics*, forthcoming.
- Barberis, Nicholas, Ming Huang, and Richard H. Thaler**, "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," *American Economic Review*, 2006, 96 (4), 1069–1090.
- Bartmann, Dieter and Martin J. Beckmann**, *Inventory Control: Models and Methods*, Springer-Verlag, 1992.
- Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein**, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, 2007, 62 (5), 2123–2167.
- Bergstresser, Daniel, John M.R. Chalmers, and Peter Tufano**, "Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry," *Review of Financial Studies*, 2008, 22 (10), 4129–4156.

- Bertaut, Carol C. and Martha Starr**, "Household Portfolios in the United States," *FEDS Working Paper No. 2000-26*, 2000.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, "Behavioral Household Finance," in Douglas B. Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics: Applications and Foundations*, Vol. 1, Elsevier, 2018, pp. 177–276.
- Black, Sandra E., Paul J. Devereux, Petter Lundborg, and Kaveh Majlesi**, "Learning to Take Risks? The Effect of Education on Risk-Taking in Financial Markets," *Review of Finance*, 2018, 22 (3), 951–975.
- Boizot, Christine, Jean-Marc Robin, and Michael Visser**, "The Demand for Food Products: An Analysis of Interpurchase Times and Purchased Quantities," *Economic Journal*, April 2001, 111 (470), 391–419.
- Bonaparte, Yosef, George M Korniotis, and Alok Kumar**, "Income Hedging and Portfolio Decisions," *Journal of Financial Economics*, 2014, 113 (2), 300–324.
- Broda, Christian and David E. Weinstein**, "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 2010, 100 (3), 691–723.
- 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 Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review*, 2015, 105 (3), 993–1029.
- Cole, Shawn, Anna Paulson, and Gauri Kartini Shastry**, "Smart Money? The Effect of Education on Financial Outcomes," *Review of Financial Studies*, 2014, 27 (7), 2022–2051.
- Davis, Steven J. and Paul. Willen**, "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice," *Quarterly Journal of Finance*, 2014, 3, 1–53.
- Einav, Liran, Ephraim Leibtag, and Aviv Nevo**, "Recording Discrepancies in Nielsen Homescan Data: Are they Present and Do They Matter?," *QME*, 2010, 8 (2), 207–239.
- Epstein, Larry G. and Martin Schneider**, "Ambiguity and Asset Markets," *Annual Review of Financial Economics*, 2010, 2 (1), 315–346.
- Fisman, Raymond and Inessa Love**, "Trade Credit, Financial Intermediary Development, and Industry Growth," *Journal of Finance*, 2003, 58 (1), 353–374.
- Food Safety and Inspection Service FoodKeeper Data**, United States Department of Agriculture 2020. Available at <https://catalog.data.gov/dataset/fsis-foodkeeper-data> (accessed July 7, 2020).
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny**, "Money Doctors," *Journal of Finance*, 2015, 70 (1), 91–114.

- Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai**, "Household Finance," *Journal of Economic Literature*, forthcoming.
- Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo**, "Consumer Shopping Behavior: How Much do Consumers Save?," *Journal of Economic Perspectives*, 2009, 23 (2), 99–120.
- Grinblatt, Mark, Matti Keloharju, and Juhani Linnainmaa**, "IQ and Stock Market Participation," *Journal of Finance*, 2011, 66 (6), 2121–2164.
- Grossman, Sanford and Guy Laroque**, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 1990, 58 (1), 25–51.
- Guiso, Luigi and Paolo Sodini**, "Household Finance: An Emerging Field," in George M. Constantinides, Milton Harris, and Rene M. Stulz, eds., *Handbook of the Economics of Finance*, Vol. 2, Elsevier, 2013, pp. 1397–1532.
- **and Tullio Jappelli**, "Awareness and Stock Market Participation," *Review of Finance*, 2005, 9 (4), 537–567.
- **, Paola Sapienza, and Luigi Zingales**, "Trusting the Stock Market," *Journal of Finance*, 2008, 63 (6), 2557–2600.
- Haliassos, Michael and Alexander Michaelides**, "Portfolio Choice and Liquidity Constraints," *International Economic Review*, 2003, 44 (1), 143–177.
- **and Carol C. Bertaut**, "Why do so Few Hold Stocks?," *Economic Journal*, 1995, 105 (432), 1110–1129.
- Heaton, John and Deborah Lucas**, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 2000, 55 (3), 1163–1198.
- Hendel, Igal and Aviv Nevo**, "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 2013, 103 (7), 2722–2751.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein**, "Social Interaction and Stock-Market Participation," *Journal of Finance*, 2004, 59 (1), 137–163.
- Hurd, Michael, Maarten Van Rooij, and Joachim Winter**, "Stock Market Expectations of Dutch Households," *Journal of Applied Econometrics*, 2011, 26 (3), 416–436.
- Jappelli, Tullio and Luigi Pistaferri**, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, 2.
- **, Luigi Guiso, and Michael Haliassos**, *Household Portfolios*, MIT Press, 2002.
- Kaplan, Greg and Guido Menzio**, "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 2016, 124 (3), 771–825.
- **and Sam Schulhofer-Wohl**, "Inflation at the Household Level," *Journal of Monetary Economics*, 2017, 91 (C), 19–38.
- Kaustia, Markku and Samuli Knüpfer**, "Peer Performance and Stock Market Entry," *Journal of Financial Economics*, 2012, 104 (2), 321–338.
- Kézdi, Gábor and Robert J. Willis**, "Stock Market Expectations and Portfolio Choice of American Households," *University of Michigan Working Paper*, 2009.

- 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.
- Linnainmaa, Juhani T., Brian Melzer, and Alessandro Previtero**, “The Misguided Beliefs of Financial Advisors,” *Journal of Finance*, forthcoming.
- 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.
- Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar**, “The Market for Financial Advice: An Audit Study,” *NBER Working Paper no. 17929*, 2012.
- 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 2004–2014. 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 2006–2014. Available at <https://www.chicagobooth.edu/research/kilts/datasets/nielsen> (accessed December 2, 2015).
- Petersen, Mitchell A. and Raghuram G. Rajan**, “Trade Credit: Theories and Evidence,” *Review of Financial Studies*, 1997, 10 (3), 661–691.
- Rooij, Maarten Van, Annamaria Lusardi, and Rob Alessie**, “Financial Literacy and Stock Market Participation,” *Journal of Financial Economics*, 2011, 101 (2), 449–472.
- Roussanov, Nikolai**, “Human Capital Investment and Portfolio Choice over the Life-Cycle,” *University of Pennsylvania Working Paper*, 2004.
- Samphantharak, Krislert and Robert M. Townsend**, *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting*, Cambridge University Press, 2010.
- Stroebel, Johannes and Joseph Vavra**, “House Prices, Local Demand, and Retail Prices,” *Journal of Political Economy*, 2019, 127 (3), 1391–1436.
- Survey of Consumer Finances**, Board of Governors of the Federal Reserve System 1983–2016. Available at <https://www.federalreserve.gov/econres/scfindex.htm> (accessed September 20, 2017).
- Vissing-Jørgensen, Annette**, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, 110 (4), 825–853.
- Yang, S. Alex and John R. Birge**, “Trade Credit, Risk Sharing, and Inventory Financing Portfolios,” *Management Science*, 2018, 64 (8), 3667–3689.

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

Online Appendix

Financial Returns to Household Inventory Management


Scott R. Baker

Stephanie Johnson

Lorenz Kueng

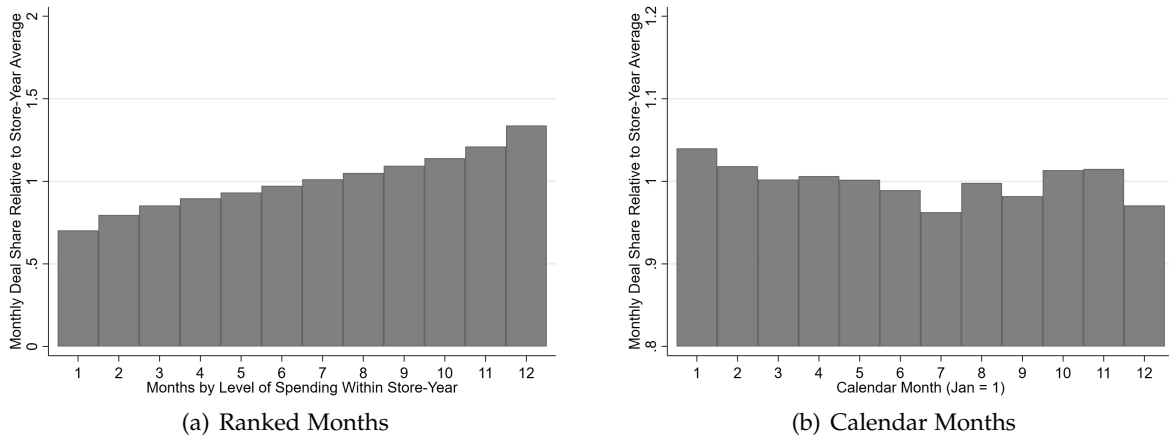
A Appendix Figures

Figure A.1 – Retailer Deal Concentration




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. A.1(a) plots the average across retailers by ranked weeks (so week 1 is the week with the lowest deal share). A.1(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 A.2 – Store Deal Concentration




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. A.2(a) plots the average across stores by ranked month (so month 1 is the month with the lowest deal share). A.2(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 A.3 – Product Life




Notes: We assign a product life measured in months to each product module and compute the share of spending on products in each monthly bin using the NCP. For food and beverage items this is based on the FSIS data. The category < 1 includes products with a life of less than 1 month; category z includes products with a life of more than z months and less than $z + 1$ months; category 12+ includes products with a life of 12 months or more.

Figure A.4 – In-store Savings and Inventory under Alternative Inventory Assumption



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 using Nielsen “departments”. 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 $\frac{\text{Annual Savings}}{\text{Annual Spending}}_i = \alpha + \beta \frac{\text{Annual Avg. Inventory}}{\text{Annual Spending}}_i + \gamma \text{Trips}_i + \epsilon_i$.

Figure A.5 – Quantity Purchased Around Move Date




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 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.

$$\log Q_{i,t} = \text{Month FE} + \text{Household FE} + \text{ZIP3 FE} + \sum_{s=-12, s \neq -6}^{12} \beta_s \text{Moved}_{i,t-s} \\ + \alpha \text{More than year prior}_{i,t} + \gamma \text{More than year after}_{i,t} + \epsilon_{i,t}$$

Where $\log Q_{i,t}$ is the log quantity purchased in ounces by household i in month t and $\text{Moved}_{i,t}$ is an indicator equal to 1 if household i moved in month t . $\text{More than year prior}_{i,t}$ is an indicator equal to 1 if household i moved more than one year after month t , and $\text{More than year after}_{i,t}$ is an indicator equal to 1 if household i moved more than one year before month t . The sample is restricted to households who moved to a new 3-digit Zip Code exactly once. Standard errors are clustered by household.

Figure A.6 – Coupon Use Around Move Date



Notes: This figure shows how the probability of using a coupon changes 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.

$$\text{Coupon}_{u,i,t} = \text{Month FE} + \text{Store} \times \text{Household FE} + \sum_{s=-12, s \neq -6}^{12} \beta_s \text{Moved}_{i,t-s} \\ + \alpha \text{More than year prior}_{i,t} + \gamma \text{More than year after}_{i,t} + \epsilon_{u,i,t}$$

Where $\text{Coupon}_{u,i,t}$ is an indicator equal to one if a coupon was used for product u purchased by household i in month t and $\text{Moved}_{i,t}$ is an indicator equal to 1 if household i moved in month t . $\text{More than year prior}_{i,t}$ is an indicator equal to 1 if household i moved more than one year after month t , and $\text{More than year after}_{i,t}$ is an indicator equal to 1 if household i had moved more than one year before month t . The sample is restricted to households who moved to a new 3-digit Zip Code exactly once. Standard errors are clustered by household.

B Appendix Tables

Table B.1 – Bulk Calibration

	(1)
$e^{-\sigma \text{Units}}$	0.88*** (0.05)
Constant	0.84*** (0.01)
Number of Observations	2,639,411

Notes: This table shows estimates of α_0 and α_1 from regression specification (18). Standard errors are clustered by product module. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table B.2 – Returns under Alternative Inventory Assumption

Working Capital Ratio	Min. Inv. (\$)	Max. Cash (\$)	% Savings: Deal Bulk		Interval Δ^* (Months)	m_0^*	Net Return (%)
0.05	132.71	102.84	9.46	11.43	0.25	4	63.03
0.10	370.19	100.55	12.66	11.52	0.25	8	12.87
0.15	596.30	104.03	13.51	11.95	0.26	11	0.81
0.20	669.33	103.98	13.68	11.95	0.26	12	0.00
0.25	669.33	103.98	13.68	11.95	0.26	12	0.00
0.30	669.33	103.98	13.68	11.95	0.26	12	0.00

Notes: This table shows results when the model is calibrated to match average inventory of \$721. This is the average inventory obtained when we use Nielsen “departments” rather than “product groups” to compute inventories. See the description in Table 6.