Debt and the Consumption Response to Household Income Shocks

Scott R. Baker
Finance Department
Kellogg School of Management*

November 2014

Abstract

This paper exploits a detailed new dataset with comprehensive panel financial information on millions of American households to investigate the interaction between household balance sheets, income, and consumption during the Great Recession. In particular, I test whether consumption among households with higher levels of debt is more sensitive to a given change in income. I match households to their employers and use shocks to these employers to derive persistent and unanticipated changes in household income. I find that highly-indebted households are more sensitive to these income fluctuations and that a one standard deviation increase in debt-to-asset ratios increases the elasticity of consumption by approximately 25%. I employ household savings and credit availability data to show that these results are driven largely by borrowing and liquidity constraints. These estimates suggest that the drop in consumption during the 2007-2009 recession was approximately 20% greater than what would have been seen with the household balance sheet positions in 1983.

*Thanks to Nicholas Bloom, Caroline Hoxby, Luigi Pistaferri, Ran Abramitzky, Steven Davis, Bob Hall, Pablo Kurlat, John Taylor, Itay Saporta, Andrey Fradkin, Pete Troyan, Siddharth Kothari, and Frederic Panier for their invaluable advice and support. Thanks also to numerous seminar participants at Stanford and the American Economic Association meetings for their helpful comments and suggestions. This research was supported by the Bradley Research Fellowship through the Stanford Institute for Economic Policy Research. The author was a paid part-time employee of the firm owning the data utilized for this paper but was not paid for work related to the paper and the paper was not subject to review by the firm prior to release.

1 Introduction

The Great Recession of 2007-2009 featured a historically large buildup in consumer debt (see Figure 1) preceding a dramatic decline in consumer spending. The possibility that higher levels of household debt induce deeper or longer recessions has important implications for both financial and mortgage regulation as well as for modeling the business cycle. More broadly, a better understanding of the dynamic relationship between a household's spending decisions, income process, and balance sheet, is imperative to accurately describing microeconomic drivers of business cycles. Despite the importance of understanding the nature of the relationship between household balance sheets and consumption behavior, the case for a causal relationship between the two has been difficult to build due to endogeneity concerns as well as limited data covering the entirety of household finances.

The fact that recessions in the United States have often been preceded by debt buildups has led to a wide range of research investigating the relationship between household debt and the business cycle.[1] A growing body of research has identified household balance sheets as important channels through which financial shocks to households can be amplified, with higher debt-to-asset ratios leading to higher implied elasticities of consumption with respect to wealth. Supporting evidence for this view can be seen in periods as diverse as recessions in the United States such as the 1990-1991 recession, the Great Recession, as well as the 'Lost Decade' in Japan.[2] In addition, theoretical work has described the pathways through which household debt could influence consumption elasticities, often by appealing to liquidity constraints or borrowing limits.[3]

However, many of the most commonly utilized incomplete-markets models take a simplified view of household balance sheets, with only a single savings instrument that abstracts from simultaneous asset and debt holdings. For instance, while a standard Bewley model predicts a decline in the elasticity of consumption with respect to income for households with higher wealth, it fails to predict any consequence of gross debt buildups or increases in debt to asset ratios conditional on net asset holdings.

This paper analyzes this relationship with new data covering a large number of American households and their employers over 6 years, from 2008-2013, and illustrates how balance sheets and access to credit shape the household consumption response to income shocks. I find strong evidence that liquid savings and access to credit drive nearly all of the heterogeneous consumption responses to income, while debt

[1] For example, see work by Glick and Lansing (2009), Isaksen et al (2011), the IMF (2012), Igan et al (2012), and Chmelar (2012)

[2] Mishkin (1977, 1978) describes both the Great Depression as well as the 1973 recession through the lens of consumer balance sheet movements. Others, including Mian and Sufi (2010, 2011, 2013), Dynan (2012), King (1994), Leamer (2007), Koo (2011), and Olney (1999) examine the path of consumption during recessions and find links to debt accumulation. Carroll and Dunn (1997) view the 1990-1991 recession as a consequence of the debt buildup of the 1980s as well as sudden worsening of unemployment rate expectations which ignited household desires for paring down debt levels. See Petev, et al (2011), French, et al (2013) and Figure B1 for discussions of the dramatic drop in income growth expectations since the Great Recession in the United States.

[3] Eggertsson and Krugman (2011) are among those to give theoretical evidence for debt-driven recessions where some agents are forced to deleverage, due to revisions in how much debt it is safe for agents to hold, thereby decreasing aggregate demand in conjunction with nominal rigidities like the zero lower bound. Guerrieri and Lorenzoni (2011), Hall (2011), and Midrigan and Philippon (2011) develop models speaking to household debt as a cause of painful recessions. As an alternative explanation, Alan, Crossley and Low (2012) find that a deleveraging and increase in the savings rate is driven by an increase in income uncertainty, modeled as an upward shift in the variance of permanent income shocks. Finally, Kreiner, et al (2013) note that, while the two are correlated, there is substantial variation in the relationship between credit and debt at a household level, emphasizing the importance of measuring each separately.

plays a minor role. This paper is the first to utilize a panel of comprehensive high-frequency household-level financial information to answer this question. This allows for more precise identification strategies, a more generalizable result, as well as the ability to highlight particular channels through which the observed effects occur, without many of the concerns that come with the more aggregated data often used in the existing literature.

An important consideration when analyzing income and consumption interactions is the fact that changes in labor market income are often endogenous, with individuals and households adjusting labor market participation at both the extensive and intensive margins in response to anticipated events. To address this endogeneity, I turn to a common source of labor market income variation, firms that employ members of the household (hereafter referred to as 'employers'). Given the importance of wage income and bonuses to household income, both positive and negative shocks to employers can drive persistent and nontrivial changes in total household income. I argue that these shocks are exogenous to any given employee and that the high-frequency nature of the data allows for explicitly testing, and supporting, the assumption that these shocks are unanticipated by the household.

In particular, I use highly detailed household financial data to provide empirical evidence that both confirms and expands standard incomplete-markets frameworks. Looking at changes in monthly and quarterly household income and consumption, I test whether consumption among households with higher levels of debt is more sensitive to a given change in income. I find substantial heterogeneity in the consumption response to income shocks across households and that this heterogeneity is systematically and causally linked to balance sheet positions.

In total, this paper makes three primary contributions. First, I develop a panel dataset featuring comprehensive and high-frequency financial data from over 150,000 households and match these households to their employers. I find that 'firm shocks' such as large and surprising earnings reports, layoff announcements, mergers, acquisitions, or write-offs significantly impact household income. As with other literature studying long term effects of labor market shocks, I present evidence that suggests these changes in household income are persistent.[4] I find less than full support for the strong permanent income hypothesis, though estimates in this paper are largely short- to medium-term and thus may understate eventual consumption adjustments.[5]

Second, I find that the elasticity of consumption with respect to income is significantly higher in households with high levels of debt. This heterogeneity persists after conditioning on household demographics or other household financial characteristics. In addition, significant effects are maintained when instrumenting for levels of debt, suggesting that the relationship between debt and consumption

[4]Davis and Von Wachter (2011) and Jacobson, LaLonde, Sullivan (1993) note that job displacement episodes have sizable and persistent effects on earnings, even when quickly transitioning to new firms. Oreopolous, Von Wachter, and Heisz (2008) find that new labor force entrants experience persistently lower wages when first entering the labor market during a recession. Roys (2011) estimates a structural model of the firm in which negative firm shocks can drive persistent declines in wages.

[5]Hall and Mishkin (1982) and Agarwal et al. (2007) find that consumers respond too greatly to temporary income shocks to be consistent with the theory. Wilcox (1989) finds increases in consumption when income increases are implemented, not when they are announced, contradicting the income-smoothing predictions of the model. In contrast, Wolpin (1982) and Souleles (1999), among others, find consumption responses and smoothing behavior entirely consistent with the theory. Aguiar and Hearst (2005) carefully document time use following retirement, highlighting the distinction between household consumption and household expenditures.

elasticities are not driven by unobservable household preferences.

Third, I show that differential household consumption elasticities with respect to income among households with varying levels of debt and debt-to-asset ratios can be explained primarily by borrowing and liquidity constraints. I find that as access to liquid assets and consumer credit increases, heterogeneity in debt ratios has little to no impact on consumption elasticities. This result is consistent with numerous studies of constrained consumers that exhibit convex elasticities of consumption when approaching borrowing limits.[6] Moreover, this finding can help to outline the mechanism by which household debt has been found to impact consumption elasticities by Mian, Sufi, and Rao (2013), Dynan (2012), and others. The strong relationship between household balance sheet positions and consumption behavior stresses the importance of more nuanced models including aspects such as endogenous and heterogeneous borrowing constraints or simultaneous asset and debt holdings. Moreover, I find that consumption elasticities are more sensitive to liquid wealth and debt levels than to illiquid wealth such as housing assets. This finding implies that simultaneous increases in mortgage debt and housing wealth, as seen during the mid-2000s, may have driven increased household consumption elasticities during the Great Recession.[7]

While I find that much of this variation is driven by borrowing and credit constraints, other channels may also play roles in influencing the relationship between household balance sheets and consumption behavior. One possibility is that household utility functions may be directly impacted by levels of debt. If households are averse to holding large amounts of debt relative to their income, a decline in income will prompt larger declines in consumption among highly indebted households in order to restore the desired debt-to-income ratio for a wide range of loss functions.

Overall, this paper's findings suggest that changes in household balance sheets, such as increases in gross debt or debt-to-asset ratios, have been important drivers of household behavior during the Great Recession and the subsequent recovery. Importantly, I provide evidence that the buildup in household debt in the years leading up to the recession increased sensitivity to declines in household income. These results point to the possibility of deeper recessions and increased macroeconomic volatility in countries with higher levels of household debt, especially when the assets purchased with debt are illiquid in nature.

The remainder of the paper is organized as follows: Section 2 describes the various data utilized in the empirical analysis. Section 3 details a number of exercises undertaken to validate various components of the data and Section 4 discusses a basic theoretical framework. Section 5 presents empirical results while Section 6 discusses a back of the envelope calculation of the aggregate effects of leverage during the Great Recession. Section 7 concludes.

[6] For one such recent study, Bishop and Park (2010) demonstrate that marginal propensities to consume drop steeply following a relaxation in binding borrowing constraints. Zeldes (1989), Johnson et al (2006), and Blundell, Pistaferri, and Preston (2008) all find higher levels of consumption sensitivity to income fluctuations among poorer and more credit constrained households. While a minority of papers, such as Shea (1995), find differently, the overall literature strongly supports the impact of credit constraints on the ability for households to smooth consumption across changes in income.

[7] Kaplan and Violante (2011) develop a model which allows for simultaneous holdings of debt and assets with the introduction of a high-return housing asset with transaction costs. Due to transactions costs making liquid and illiquid wealth imperfect substitutes, they use this model to explain why even wealthy households with substantial net assets may behave as though borrowing-constrained.

2 Data

2.1 Household Financial Data

The household financial data used in this paper comes from a large online personal finance website. The site provides a service that connects users' financial accounts so that user can see all of their accounts in a single location. The site allows for users to easily see summaries of their income, spending, debt, and investments across all of their accounts and has other features such as budgeting or financial goal-setting. The site has grown rapidly, from under 300,000 users in 2007 to more than 3 million active users by 2012. This large user-base has yielded a database of more than 5 billion transactions across over 10 million individual accounts. These accounts span all manner of household financial products including checking accounts, savings accounts, credit cards, loans, property and mortgage accounts, equity portfolios, and retirement accounts.

The basic data elements are individual bank and credit card transactions that are automatically recorded through links to users' banks and credit providers. Each transaction is time-stamped and has information about the other party, whether it was a credit or debit, and contains a full description of the transaction as you would see on your monthly bank or credit card statement. From this merchant and descriptive data, the site automatically categorizes each transaction into one of over 100 categories (such as 'Groceries', 'Gasoline', 'Student Loans', 'Fast Food', or 'Pet Stores') in order to provide easily readable spending and income breakdowns to the user. From these data, I derive measures of total household spending and income as well as subsets of income and spending based on the categorization of the transactions. This categorized spending and income data provides an invaluable source of information regarding each individual user's financial means and behavior.

In addition to transactional data from bank and credit card accounts, the site automatically archives daily balance data for equity, retirement, property, real estate, and loan accounts. The loan and property data, in particular, form the basis of important components of analysis in this paper. Loan data consists primarily of car loans and home mortgages, while property accounts denote the value of home prices over time for each homeowner. I use this data to determine gross and net housing wealth at a household level. Often accompanying this balance data is information on interest rates, types of accounts (such as IRA or 401k), and the bank or firm providing the account.

One major advantage of this data source over other sources, such as government survey data like the Current Population Survey or the Consumer Expenditures Survey, is that the data are obtained directly, automatically, and continuously from the households' financial institutions. This greatly minimizes measurement error and biases in recollection about both financial flows as well as stocks (in terms of both errors in amounts and errors in timing of consumption or income) that are endemic to survey data. A second advantage is that it provides comprehensive income and spending data for all households in the sample, avoiding potential external validity problems when compared to other datasets that solely include information on, for example, household spending on food. Such narrower datasets may result in biased estimates when households of different characteristics (e.g. income level) may systematically differ in the proportion of their budgets that are spent on food, certain types of durables, or other particular categories of spending. For both of these reasons, this data provides a more comprehensive and unbiased source of household financial data than other previously used sources.

In addition, household data allows for a much wider range of financial situations relative to geographically aggregated data. For instance, across counties in 2006, average debt-to-income ratios ranged from approximately 1 to 3 and average mortgage-to-house value ratios ranged from about 0.30 to 0.95. In contrast, individual households naturally spanned a wider range of balance sheet positions, with debt-to-income ratios of 0 and in excess of 15. Similarly, mortgage-to-home value ratios ranged from 0 to over 2, with many households underwater on their homes and thus having a ratio of more than 1. Allowing for the consideration of these more extreme balance sheet positions is crucial for accurately characterizing the microeconomic drivers of the 2007-2009 recession.

An important consideration when dealing with this data is to consider the number and type of linked accounts for a given user over time. As the site was rapidly expanding in recent with a multitude of new users who were gradually linking financial accounts, looking at simply the per-user spending patterns could potentially give a distorted view. If an individual had two credit cards and had only linked a single one, when they add the second card the site would register a large increase in spending whereas there was not truly a real spending increase. To combat this bias, I employ three cleaning and robustness strategies. First, I exclude data from a user's first 3 months using the site, as this period is typically when most of the adding of accounts takes place. Second, I exclude users with highly volatile numbers of accounts or insufficient numbers of accounts (require ≥ 3 accounts), as it seems likely that there is unobserved actions being taken here that I cannot control for adequately. In addition, I also perform robustness tests for most specifications where I utilize spending or income per account instead of total spending or income across accounts as an alternate measure. I find results that are qualitatively unchanged, suggesting that the mitigating steps taken to address this potential bias achieved their aim.

In addition, I test whether demographics or finances affect how often individuals utilize the website. This could be an issue if certain types of users utilize the site more frequently and thus any changes to their financial circumstances are made more salient. I find that younger individuals and higher income individuals are weakly more likely to check the website. However, these effects are not dramatic, with a one standard deviation in age or income only yielding a 32% and 14% change in site visits per quarter, respectively.

Another concern could be related to which individuals a given user account covers. If a married couple possess two separate accounts on the website with some overlapping joint accounts, I may have only a distorted view of each 'user's' finances. Thankfully, in conjunction with the financial data, users provide demographic information such as age, sex, marital status, and the size of the household. Users also list whether they are a homeowner, their profession, their level of education, their income level, and their location. Using this data, I test how usage patterns relate to self-reported demographic characteristics. Testing whether I can identify the same transactions occurring on multiple users' accounts, I find few instances of these overlapping transactions, suggesting that few users have a joint account listed that is also listed by another user.

In addition, I find that users who self-report being married are dramatically more likely to have joint bank accounts, multiple savings accounts, multiple paycheck streams and users with children can be observed spending significant amounts on child-related activities and education. Overall, site usage patterns point to the conclusion that linked financial accounts cover the entirety of a household. Thus, I equate a user of this financial website with a head-of-household in the Current Population Survey (CPS) or a 'consumption unit' in the Consumer Expenditure Survey (CES).[8] For example, a 'user' represents the entirety of household spending for married couples but only represents an individual's spending for an unmarried individual living with roommates.

Being a software start-up, the demographics of the website were initially very different than those of the national as a whole. Key user characteristics like gender and age were starkly different than the national distribution in 2007 (being younger and more male). The divergence from national averages presents both benefits and problems. Though the unweighted data cannot be directly used to obtain unbiased estimates that can be extrapolated to the entirety of the US population, I am able to capture much of the spending by high-income earners that is often missed in other survey-based datasets like the Consumer Expenditure Survey that under-sample rich households.

While the demographics of the user-base were initially very different, they have become much closer to a representative national distribution by 2013 as the user-base grew dramatically. Moreover, conditional on observable household demographic and locational characteristics, financial behavior among the users seems to track closely to national averages. Validation of the ability to transform the observed financial data series into a relevant and representative sample of American households using CPS household weights is covered extensively in Section 3. CPS-weighted and unweighted summary statistics of the sample population can be found in Table 1 and Table B1.

One drawback of the data is that it does not have as complete a coverage of cash or check transactions relative to credit and debit transactions. Cash transactions can only be fully observed when a user manually enters them, though inferences about other cash transactions can be made as cash withdrawals from banks and ATMs can be directly observed. An estimated 6-8% of total spending is done with cash in the United States, compared to approximately 3-4% of spending done with cash in the sample data. A larger omission could potentially stem from check transactions, which can be observed but cannot be automatically categorized. Such transactions make up about 15% of total observed spending among households in the sample. Through inspection of manually re-categorized checks (check transactions where the user manually changed the category and description of transaction type), I find that check transactions are primarily driven by large, regular transactions like mortgage payments, rent, and utility bills. In my desired sample of employees of publicly-listed firms, paychecks make up a very small portion of these check transactions, as the vast majority of public firms' employees are paid by direct deposit (see NACHA 2010 PayitGreen survey), which can be almost perfectly observed with their correct category coding.

A second drawback is that I cannot observe withdrawals from an individuals' paycheck prior to their receiving it. For example, transactions like healthcare premiums, 401k contributions, and FICA payments are often unobserved (if I do not see the retirement account itself), as I only see the resultant paycheck that is deposited into a checking account, not the originating payment. Thus, for all households drawing a paycheck, I generally measure post-tax and post-benefit (401k, healthcare) pay, thereby

[8] Demographic information for a user only encompasses a single person, while the financial information seems to cover all household members. For this reason, I take the self-reported demographic information on the user's profile to be that of the 'head-of-household' for the purposes of CPS comparison.

estimating based on take-home income rather than total gross pre-tax income.

2.2 Firm Shocks Measurement

One potential measurement issue with looking at changes in household income is that households may adjust current spending in anticipation of future changes in household income that are unobserved by the researcher. To ameliorate this issue, I turn to using shocks originating from household employers to instrument for changes in household income. I largely employ firm shocks rather than firm stock returns primarily to help minimize the endogeneity and persistence in stock returns at a high-frequency level, picking out singularly large shocks to the firm. These shocks have sizeable impacts of household income and, in addition, there is reason to believe that the exclusion restriction holds and that shocks to the firm impact household spending solely in ways that are captured by household income. While it is possible that firm performance may influence future career prospects and thus current spending patterns, it seems likely that brighter career prospects would be accompanied by higher contemporaneous pay.

I utilize three sources to compile data regarding large shocks to firms, a summary of which can be found in Table 1. The first source data from the SEC's required 8-K filings for public firms. These filings are required by firms for a wide range (over 20 types of events) of firm-specific actions such as firm or plant closures, changes to the board or principle officer, or significant merger and acquisition activity. These filings were instituted as an effort to increase mandatory disclosure of pertinent information for shareholders. Moreover, the filings must be promptly released after any such event, with the SEC requiring release within the four days following the event itself.

For the purposes of this project, I focus on firm layoffs, acquisitions and sales completions, large write-offs, and large earnings report surprises. These shocks share three beneficial characteristics. First, they are relatively clear in their interpretation and scope. Second, they have sizable firm equity movements associated with their occurrence and thus are plausible drivers of employee earnings and earnings expectations. Finally, both investors and employees do not seem to exhibit any foresight regarding the announcements and subsequent equity movements, with little change in stock prices or consumption among employees prior to the shock. Timelines of these events are shown in Figure 2.

The second source for firm shocks is the Institutional Brokers' Estimate System (I/B/E/S) which enables me to collect data on quarterly earnings report for all publicly listed firms in my sample. From this data, I construct indicators for particularly surprising positive or negative earnings report, where a large positive earnings report surprise is:

$$\frac{(EPS - E[EPS])}{SharePrice} > 0.02$$

and a large negative earnings report surprise is:

$$\frac{(EPS - E[EPS])}{SharePrice} < -0.02$$

That is, I categorize an earnings report as a 'large earnings surprise' when the difference between earnings per share and expected earnings per share is larger than 2% of the firm's share price. A graph of large positive and negative earnings report surprises over time is displayed in Figure 2. A 2% threshold was used to approximately capture the top and bottom 1% of earnings announcements. Results are robust to using cutoffs of 1% or 3%.

Finally, I use a news-based strategy of identifying layoffs in order to provide an alternate measure of firm layoffs that also gives an intensity scale (rather than simply an indicator of layoffs). I compile a database of firm layoffs through the use of the Access World News Newspaper database. When reporting on layoffs, newspapers tend to construct titles using a set format; a practice confirmed through manual inspection of a large number of known layoffs, extended trials of search terms, and through talking to business journalists. With this in mind, I use this database of over 1,500 US newspapers to search the archive for article titles that mention each firm's name, or common shorthand for the firm's name, as well as a set of terms indicative of layoffs.[9] Thus, this query compiles a set of all articles that have titles similar to "Wells Fargo to cut 4,000 jobs" or "Alcoa lays off 2,400 workers".

This search returns approximately 5,000 articles since 2008 covering almost 400 firms. I attempt to exclude false positive matches by removing matches with fewer than 3 articles about a layoff on a given day. Furthermore, given the structure of the title of these articles, I am able to extract the number of individuals laid off during each episode and validate this number by checking the matched number of layoffs across articles regarding the same firm and day. The final sample covers 113 firms and 246 separate instances of layoffs in the sample period. A sample of these layoffs are reported in Appendix Table B2.

Table 2 tests for anticipation of these shocks to firms among both employees and investors. Columns (1) - (4) regress firm returns and a selection of firm shocks on lags of various pieces of observable household financial information such as income, wealth, and spending. I run these regressions at a high-frequency, monthly, level and find no evidence for changes in income, spending, or wealth among employees in the months leading up to layoffs or large positive or negative earnings announcements. Column (5) demonstrates the same principle at a firm level, with stock returns not reacting to shocks in the months before the shock occurs, suggesting no anticipation of the shocks by investors. This contrasts with a strong response in the stock price seen after the shocks take place, as seen in column (6). This paper currently limits the effect of the various firm shocks to shifts in the level (or growth rate) of employee earnings and assumes that this mean effect is thereafter anticipated by the affected employee. I do not explicitly consider changes in the variance of employee earnings over time based on these firm shocks.

In addition to being seemingly unanticipated, I find that these shocks to employers exhibit a high degree of persistence in terms of their effects on household income. Figure 3 plots logged average household income levels before and after positive and negative shocks to household employers after removing a time trend. I find persistent effects on household income remain up to 6 quarters following both positive and negative shocks.

Because my data is limited to a 5-year period, I cannot reject that the impact of a given firm shock dissipates gradually over the ensuing years. However, highly persistent effects on income would be consistent with much prior work finding that shocks to firms, job displacement episodes, and having

[9] Articles must contain a term from the set (positions, jobs, employees, workers, notices, roles, staff, personnel) as well as a term from the set (slash, slashes, slashing, lost, losses, layoffs, sheds, axes, cuts, fires, layoff, shed, axe, cut, cutting, axing, shedding, reduce, "lay off")

bad initial labor force experiences can affect long-term wages.[10] In general, household income derived from employers is highly persistent. When estimating the wage process as a quarterly AR1, I find a coefficient on lagged firm-related income of over 0.96. Total household income is somewhat less persistent, with an estimated coefficient on lagged total household income of approximately 0.92.[11] This estimate is consistent with other estimates of the household income process which often find AR(1) coefficients in excess of 0.90.

2.3 Firm Matching, Firm Shock Power, and User Selection

I match users to their employers using textual descriptions from users' direct deposit transactions. Direct deposit transaction descriptions are generally characterized by indicators that the transaction is a direct deposit, a string representing a firm, and anonymized identifiers.[12] Using a flexible natural language processing algorithm, I compare these paycheck descriptions to a seed list of firms to match to. The pool of firms used is the universe of publicly listed firms in both the NASDAQ and NYSE. For each firm, the matching algorithm is allowed to ignore things like punctuation such as hyphens, periods, or apostrophes and allows for abbreviations of the full firm name. In addition, I include a large number of unconventional abbreviations obtained through manual inspection of the largest 100 firms by revenue as well as the largest 100 firms by employment in the sample (e.g. 'TGT' is an identifier for a direct deposit from Target). Looking across all users of the website, the set of users able to be matched to their employers contains 1948 employers employing over 700,000 household members. The set covers employers from all major sectors, including retail, technology, banking, industrial/manufacturing, media, and professional services. Many matched firms employ large numbers of users, with several firms employing over 5,000 users each.

I limit my analysis to a sub-sample of users able to be linked to a publicly listed employer for a period of at least 4 quarters. Moreover, I exclude users matched to firms that have fewer than 50 observed users in my data, ensuring at least a moderate level of within-firm variation across individuals and making it less likely that a firm was matched to a user erroneously through a rare or custom transaction description.

One method in which I can jointly test whether the firm-employee matching algorithm was accurate and whether employees react to their employer's performance can be seen in Appendix Table B4. This table displays some quantification of the relationship between employer performance and consumer spending. Here, I match individuals in my sample to their employers during the duration of their tenure there. In addition, I construct a variable matching them to the same employer for up to 6 months prior to their working there and up to 6 months afterward if they leave.

For instance, if I observe an individual working for General Electric from July 2010 to September

[10]Davis and Von Wachter (2011) and Jacobson, LaLonde, Sullivan (1993) note that job displacement episodes have sizable and persistent effects on earnings, even when quickly transitioning to new firms. Oreopolous, Von Wachter, and Heisz (2008) find that new labor force entrants experience persistently lower wages when first entering the labor market during a recession. Roys (2011) estimates a structural model of the firm in which negative firm shocks can drive persistent declines in wages.

[11]Both estimates remove quarter fixed effects to control for seasonality.

[12]Some examples of such descriptions are: "21ST CENTURY INS DIR DEP", "SAVINGS DEPOSIT - DIRECT DEPOSIT FROM TGT PAY REG SALARY", and "GOOGLE, INC DIRECT DEP"

2012, I match his consumption behavior to the stock returns for General Electric from January 2010 to December 2012, including the 6 months before he began working for General Electric and 3 months after he left. If employees do respond to their employers' performance, captured by their stock price, we would expect to find that household spending responds to stock prices during the period of employment with the firm, but not before or afterward. Columns (1) and (2) show results of simply regressing household spending on employer stock returns for the period before, during, and after employment at a given firm. We find mild positive impacts of stock returns on spending, though insignificant without financial controls. Columns (3) and (4) give results from a regression of household consumption on employer returns (inclusive of the 6 months preceding and following being matched to the employer) as well as an interaction of employer returns with an indicator for whether the individual currently works at that firm. Here we see that the interaction term's coefficient is positive and significant while the uninteracted term is insignificant. This suggests that individuals are reacting to current employer stock price and that the firm-employee matching algorithm performed reasonably well given that they do not react to their future employers' stock price prior to being hired or following a departure.

Moreover, a household's employer has little predictive power for household debt-to-asset ratios, with each firm employing a range of employees with widely varying income, wealth, and debt characteristics (see Figure 4 for distributions within several selected firms). Running a regression of 2008 household debt-to-asset ratios on household demographic and locational characteristics, as well as employer fixed effects, I find unadjusted significant effects for only 5 out of more than 200 employers in the sample. Households may vary in their financial characteristics, even controlling for employer, demographics, and income, for a variety of reasons. For instance, households may have had varying shocks in the past, such as medical incidents, large inheritances, or differing marital circumstances. There may exist heterogeneity among otherwise similar households in the expectations about future financial paths or in inherent traits such as risk aversion or discount rates. In addition, households may vary in the degree to which other behavioral factors such as "mental accounting" or "rule of thumb" behavior influences spending and asset accumulation decisions. Finally, household characteristics do not seem to predict being employed by a firm that is subject to more or fewer firm-level shocks. Running a probit of firm shocks on household demographic and locational characteristics, I find no significant effects.

Shocks to a household's employer have moderately large impacts on subsequent household income. In Table 3, I test whether firm-driven shocks impact income of various types among employee households. I consistently find negative effects on household income following 'negative' shocks such as large negative earnings surprises, layoffs, and material impairments while finding positive effects of large positive earnings surprises and completions of acquisitions or dispositions. For instance, from column (2) I find large positive earnings surprises tend to increase overall income by about 0.7-1.1%. Similarly, large negative earnings surprises decrease spending by about 1.3%, Column (3) shows that non-firm based income (such as rental income, stock dividends, or spousal income from a different firm) is unaffected by shocks to a household's employer. This is evidence that there is little offsetting income response to firm shocks along other margins and that these shocks do represent changes in household income and not just a shift of income from one source to another.

Columns (4) and (5) show the impact of these firm shocks on paycheck income separately from bonus income, which displays much higher effects given the more dramatic volatility inherent in bonuses. Importantly, paycheck income, a more persistent component of employee compensation, is still significantly affected and points to persistent effects on household income.[13]

Column (6) collects the firm shocks into two bins, with large positive earnings surprises and acquisition/disposition completions defined as positive shocks and large negative earnings surprises, layoff announcements, and material impairments defined as negative shocks. I again find positive effects of positive shocks on both spending and income and negative effects of negative shocks, with magnitudes of the order of 0.8%-1.2% increases or declines in income and spending following a shock. Overall, shocks to employers are strong drivers of changes in paychecks and bonuses in particular and household income in general. Moreover, the shocks are moderately persistent, with lower levels of income remaining one year following the various firm shocks.

These results suggest the potential for additional work examining the mechanisms behind household income adjustments driven by firm shocks. Table 3 finds significant effects both on regular income as well as bonus income at a fairly high-frequency level. Splitting the two sources of the decline, I find that the impact on bonus income occurs more rapidly while declining over time while the impact on payroll income grows and stabilizes over several quarters. However, some of the observed impact on regular payroll income may in fact be bonus income rolled into regular paychecks. In addition, it may be the case that hours adjust for some employees following good or bad news for the firm. For instance, a firm reporting extremely bad earnings numbers may be under pressure to cut costs and cut hours assigned to employees. Unfortunately, the data is unable to distinguish between changes in hours and changes in rates of pay and may conflate the two.[14]

3 Data Validation

One concern is whether users have linked sufficient accounts to the site to get an accurate picture of their finances. A random survey of 3,649 site users in 2011 provides some reassurance on this point. It found that over 95% of users had linked all or almost all of their checking accounts and over 93% of users had linked all or almost all of their savings accounts. In addition, 91% of users had linked all or almost all of their credit cards. While these accounts had the highest coverage rate, over 75% of users had linked all or almost all of their equity accounts. Similarly, over 90% of users who self-identified as homeowners included information about their home value or mortgage. This evidence points to near-complete coverage of active users' accounts that get the most frequent use and attention. Despite this, we find less complete coverage of more unconventional asset accounts such as retirement accounts, with only approximately 50% of users linking all or almost all of this type of account. This pattern seems to derive from the fact that the most common use of the website is for tracking income and expenditures so

[13] Non-parametric time series analysis finds stronger impacts on bonus income in the near-term (under two quarters), while changes to paycheck income occur over a generally longer period (1-3 quarters)

[14] Appendix Table B6 presents results from specifications mirroring the main results in Table 4 but utilizing solely changes in logged paycheck income or bonus income as a dependent variable. I find that spending is more elastic in response to changes in paycheck income than overall household income. Moreover, looking only at bonus income, I find lower levels of elasticity and no significant association with levels of household debt, consistent with the fact that changes in bonus income are fairly transitory.

categories that see the most frequent transactions are those that are most likely to be linked. Moreover, when asked why they had not linked all of their financial accounts to the site, the most common answer given was that there was little activity on the accounts or that they felt no need to track those accounts. This fact reinforces the view that the unlinked accounts most likely do not account for an appreciable amount of transactions or activity relative to the linked accounts.

While the demographics of the user-base of the site have become more similar to those of the nation in recent years, a worry about representativeness still exists. To combat this worry, I benchmark the aggregate data against several known measures of national consumer behavior and use CPS-derived weights to synthetically approximate the distribution of households in the United States.

First, I compare observed spending among users to monthly data from the Census Retail Sales numbers. Census Retail Sales data comes from a monthly survey of 3,000 large retailers and over 9,000 small retailers and data are broken down by type of retailer. I match my observed categorical spending data to the relevant categories and construct average per-user categorical monthly retail spending, weighting users by CPS weights for age, sex, income, and state of residence. Figure 5 shows results for a range of categories like gasoline sales, electronics, restaurants, and clothing sales.

It is important to note divergences from the site's data from some categories of Census Retail spending. The motor vehicle spending panel of Figure 5 shows data for automobile and motor vehicle sales, which do not track well between the two sources. This is primarily because my data observes payments from a consumer's side, while Census data measures spending from the retailer's side. For some categories like motor vehicles, these two aspects diverge as many individuals purchase cars with loans or other financing and pay them off over time. Thus, I see more gradual changes in motor vehicle spending as I often observe monthly payments rather than the lump sum purchases of vehicles that dealerships and retailers record. This phenomenon is also present, though less pronounced, for other durables categories in which consumer credit or financing is often used.

For non-motor vehicle categories of spending, the average correlation between my observed spending data and the Census Retail data is 0.86, with this average correlation rising to 0.89 when also excluding furniture sales which are also commonly bought on credit. A number of categories, such as Gasoline, Clothing, and Restaurants, have monthly correlations with the Census Retail data well in excess of 0.90. Finally, the data from the site compares well to annualized data from the Consumer Expenditures Survey, with a correlation of over 0.87 across all overlapping categories of spending and time periods (2009-2013).

These validation exercises also serve to further alleviate some worries about the completeness of the accounts. If there were systematic biases for individuals in terms of the accounts that they chose to link to the site, overall trends in spending and income would be unlikely to so closely match national trends. The fact that this holds true for so many categories of spending cements this view.

Individuals in the sample, when weighted by observables, also match national distributions of things like housing wealth. Figure 7 shows in-sample logged average reported home price by year and zip code compared with the overall logged average home price by year and zip code provided by Zillow.com. The two have a correlation of greater than 0.88, showing a strong relationship between reported and observed house prices. While the unweighted sample displays systematically higher house prices and housing wealth than the overall averages, when weighted by observables, the sample corresponds much more closely to the national distribution. I also test observed distributions of liquid savings and credit among households in my data against the 2010 Survey of Consumer Finance, finding close matches (see also Figure 8, debt ratio figures are also comparable to SCF distributions). These results suggest that not only do weighted spending trends match national distributions, but observed asset and debt levels are consistent, as well.

As a result of these validation exercises, all regression results in this paper are weighted on (head-of) household observables such as age, sex, income range, and state of residence according to CPS weights, yielding results with a high level of external validity when extrapolating to the nation as a whole.[15] Weights are recalculated each year in the sample period (2008-2013).

The final sample is restricted in a number of ways to maximize statistical power as well as minimize measurement error. Users must also have more than 3 linked financial accounts (e.g. savings, credit card, or loan accounts) and have filled out demographic information about age, sex, marital status, and location in order to ensure they are active users with complete account profiles. I exclude users with more than 25% of their spending deemed as uncategorized or who have persistent transfers of funds to unlinked accounts (eg. credit card payments to an active credit card that has not been linked). I also require users to be able to be matched to an employer for at least 12 months and to have logged into their accounts sometime in the last 6 months of 2012. The latter qualification helps to exclude inactive users who's financial information and account linkages are likely to be more out of date. Finally, I exclude users with large internal discrepancies such as self-identified homeowners who do not enter any information regarding their home values or users who list income that is 50% less than or 100% more than their observed level of income. The final sample includes 156,604 households.

4 Theoretical Framework

As an illustration, I turn to the much-studied incomplete markets Bewley model for insight into a framework with heterogeneous shocks and household savings as a consumption smoothing device. In a simple version of this model, agents earn an exogenous stream of wages and have access to savings bonds as their only form of insurance.

Finitely lived households maximize utility:

$$\sum_{t=1}^{T} \beta^{t-1} U(c_t) \text{ where } U(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}$$

with a budget constraint $b_{t+1} = (1+r)b_t + y_t - c_t$ and a borrowing limit of $b_t \geq \Theta\ \forall t$. Household wages follow an AR1 process: $y_t = \alpha + \rho y_{t-1} + \epsilon$ with $\epsilon \sim N(0, \sigma)$

This framework results in a standard Euler equation of the form:

$$u'(c_i, t) = E_t[(\beta(1+r_t)u'(c_{i,t+1})]$$

[15] Results are qualitatively robust to performing regressions without using CPS household weights or using weights based on the distribution of private sector employees rather than total household numbers.

when borrowing constraints are slack. If constraints bind, future consumption will be necessarily more responsive to a given change in income.

I solve this model quarterly with standard parameter assumptions.[16] I calibrate the income and asset distribution to match the observed distribution in the data based on the first and second moments of the data and map the wage process to a discretized 11-point Markov process. In addition, I utilize a high level of persistence for the wage process, with $\rho = 0.9$, which is approximately equal to that seen in the overall income process in the data. Results are qualitatively similar when varying ρ within reasonable limits.

I compare results from this model (I run the model for 5 years for burn-in and then take the next 5 years of data to analyze) to equivalent data from my panel of users. For both datasets, I estimate regressions examining how changes in income and net asset holdings affect household consumption in Table 4, columns (7) and (8). At any point in time, the number of remaining periods, shocks to household income, and the net asset position of the household affect consumption behavior. I estimate regressions of the form:

$$\Delta ln(c_{it}) = \alpha \Delta ln(y_{i,0}) + \beta \Delta ln(y_{it}) * w_{i,0} + \lambda h_i + \eta_{it}$$

where h represents household fixed effects, proxying for periods left in a household's lifetime and w is equal to a household's net assets. I employ household fixed effects rather than age as I cannot precisely determine likely times of retirement or death and cannot otherwise account for the timing of other lifecycle events that drive short to medium-run consumption trends. I include seasonal effects in both the data and the model analysis though seasonal effects have little impact in the model as there are no explicit holidays or yearly trends. Net asset data is calculated at the beginning of the sample period and held constant over time.

While being a useful tool for economists, basic Bewley models and other models with single types of savings mechanisms does not predict any effects of increasing both gross debt and gross assets, as households never hold both assets and debt at the same time. For instance, if a household increases its gross debt and gross assets by \$1 each, net worth remains the same and thus behavior would be predicted to be unchanged. I discuss empirical results that expand on this simple framework in the following section. In addition, given the lack of long-term time-series data for my sample households, I primarily turn to the exogeneity of the income shocks that I consider rather than relying on time-series properties of the Euler equation and income process to recover causal relationships.[17]

In the following section, I attempt to enrich this framework by looking at various measures of indebtedness, by separating household balance sheet positions, and also by including additional information on liquidity and credit constraints that may be a more direct driver of household consumption behavior. In addition, I attempt to focus solely on shocks that drive more permanent changes in income (noted by σ):

$$\Delta ln(c_{it}) = h_i + \alpha \Delta ln(y(\sigma)_{it}) + \beta \Delta ln(y(\sigma)_{it}) * Leverage_{i,0} + Season_t + \eta_{it}$$

[16] Parameter choices include $\gamma = 2$, annualized $\beta = 0.95$, the real interest rate is fixed at $\frac{1}{\beta} - 1$.

[17] As in Johnson, et al (2011), who observe the impact of randomly timed tax rebates on household consumption free of Euler equation assumptions.

$$\Delta ln(c_{it}) = h_i + \alpha \Delta ln(y(\sigma)_{it}) + \beta \Delta ln(y(\sigma)_{it}) * Debt_{i,0} + \beta_2 \Delta ln(y(\sigma)_{it}) * Assets_{i,0} + t_t + \eta_{it}$$
$$\Delta ln(c_{it}) = h_i + \alpha \Delta ln(y(\sigma)_{it}) + \beta \Delta ln(y(\sigma)_{it}) * Leverage_{i,0} + \gamma \Delta ln(y(\sigma)_{it}) * Credit_{i,0} + Season_t + \eta_{it}$$

5 Results

5.1 Income Shocks and Debt

In Table 4, I investigate the impact of debt in amplifying the consumption response to household income during the 2008-2013 period. I use a detailed quarterly panel to look closely at the dynamic relationship between balance sheets, income, and consumption. I employ a panel fixed effects strategy as well as an instrumental variables approach in order to isolate the causal impact of debt on the sensitivity of household consumption to income. As noted above, in all specifications run at a household level, I weight users by various demographic and locational observable factors (age, sex, income bracket, state of residence).[18]

Column (1) determines the baseline elasticity of consumption with respect to income in a panel fixed-effects OLS framework. In column (2), I use the same approach to examine the effects of debt to asset or debt to income (collectively labeled 'leverage' in specifications below) ratios on household consumption behavior. For all specifications, I denominators of 'leverage', household assets or household income, are derived from 2008 values such that changes in household income or assets over time do not influence measured 'leverage':

$$\Delta Spending_{it} = \beta_0 + \beta_1 \Delta Income_{it} + \beta_2 \Delta Income_{it} * Leverage_{it} + \beta_3 Leverage_{it} + \beta_i HH_i + \beta_t Time_t + u_{it}$$

As noted in Kaplan and Violante (2010) and others, time until retirement or death often has a profound and systematic effect on consumption patterns. In lifecycle models with either known or uncertain ages of retirement and death, transaction costs of shifting assets between liquid (such as savings, checking, and brokerage accounts) and illiquid holdings (such as real estate holdings and retirement accounts) can produce sizable periods of trending consumption, even without any coincident changes in income. Given the relatively short time-span of household behavior that I observe, I cannot fully model household lifecycle decision-making, and my estimates would be influenced by the stage of life that households were currently in. Thus, I attempt to control for these short-term lifecycle trends through the use of household fixed effects. These allow me to observe deviations from consumption trends at a household level that respond to income and balance sheet effects rather than lifecycle pressures.[19]

In this specification, spending is measured as total logged quarterly household spending and income measured total logged quarterly household income. I interact the change in logged quarterly household income with the ratio of total debt to total assets and examine how differences in leverage are associated with differences in the responsiveness of household spending to income. I find that higher levels of

[18] All regressions employing income or spending variables will be utilizing natural logs of income and spending. All household panel specifications also include household fixed effects to control for systematic trends in household spending patterns. Unless otherwise noted, standard errors are clustered at a household level.

[19] The empirical results are qualitatively similar when utilizing age dummies instead of household fixed effects, with strong impacts of age on the trends in consumption over the 4-5 years that I observe households.

leverage are significantly related to higher sensitivity of household spending to income. For instance, an increase of one in the debt-to-asset ratio (about 1.25 of a standard deviation movement) sees an increase in spending by an additional 20% relative to the baseline change in spending for a given change in income. Note that all specifications include household fixed effects, so estimates of income and spending changes are based on deviations from household-level trends. Households generally exhibit strong trends over time as they age and obtain higher-paying jobs and raises within jobs. Estimates including only time fixed effects or time fixed effects in combination with demographic characteristics of households perform qualitatively similarly.[20]

Column (3) performs the same procedure but substituting debt to income ratios for debt to asset ratios. I find significant and positive effects of debt to income ratios on consumption elasticities.[21] These results point to a large degree of heterogeneity in consumption elasticities across households with differing balance sheet positions.

Instrumental variable results are also shown in Table 4, in columns (4)-(6). In these columns, I employ instruments in order to isolate exogenous levels of both debt and changes in income. I utilize positive and negative firm shocks to instrument for changes in household income. In addition, I include a widely used instrument for household debt in 2008, housing supply elasticity, that I discuss in detail in the Data Appendix.[22] In addition, I interact this variable with each household's homeowner status in 2008 to provide additional variation in household debt accumulation paths. In each case, household and time fixed effects are included:

$$\Delta Spending_{it} = \beta_0 + \beta_1 \Delta \widetilde{Income}_{it} + \beta_2 \Delta \widetilde{Income_i * Leverage_i} + \beta_i HH_i + \beta_t Time_t + \xi_{it}$$

Where the variables $\widetilde{\Delta} Income_{it}$ and $\Delta \widetilde{Income_i * Leverage_i}$ represent fitted values from the following regressions:[23]

$$\Delta Income_{it} = \beta_0 + \beta_1 FirmShockPos_{it} + \beta_2 FirmShockNeg_{it} + \beta_3 Saiz_i * FirmShockPos_{it} + \beta_4 Saiz_i * FirmShockNeg_{it} + \beta_i HH_i + \beta_t Time_t + \xi_{it}$$

$$\Delta Income_{it} * Leverage_i = \beta_0 + \beta_1 FirmShockPos_{it} + \beta_2 FirmShockNeg_{it} + \beta_3 Saiz_i * FirmShockPos_{it} + \beta_4 Saiz_i * FirmShockNeg_{it} + \beta_i HH_i + \beta_t Time_t + \xi_{it}$$

Column (5) gives results for the 2nd stage of this regression, finding stronger effects of income changes and income changes interacted with debt-to-asset ratios relative to column (2). I find that both

[20] One potential concern may be that individuals switch jobs in response to shocks to their firms. I discuss the observable incidence of job-switching in the appendix, finding little to suggest that this is a major concern. The changes in income here are generally not very large, making it less likely that individuals would incur search and switching costs to alleviate income declines.

[21] The correlation between the Debt-to-Asset ratio and Debt-to-income is 0.669. Correlation taken across 156,604 households in the sample.

[22] This measure was developed by Albert Saiz (2010) and is derived from geographic differences in the potential elasticity of housing supply across MSAs.

[23] Also included in the first stage regressions are triple-interactions with the homeowner status of households, included to increase first stage power, as the Saiz instrument primarily affects homeowners through property price appreciation. All standard errors corrected for 2SLS. Results all pass Sargan overidentification test.

the coefficient on the change in logged income as well as the coefficient on the interaction term increase. This increase in measured responsiveness may be partially due to the fact that the changes in income driven by firm shocks are more unexpected than typical changes in quarterly income that households may anticipate and plan for, thus such shocks produce a stronger short-term consumption response. Again, spending among highly indebted households is significantly more responsive to changes in income relative to low-debt households, both economically and statistically. A one standard deviation increase in debt-to-assets would again raise the elasticity of household spending to household income by about 20% (by about 0.07 from a baseline of 0.37, given a median debt to asset ratio of approximately 0.4). Column (6) utilizes the gross debt-to-income ratio as an alternate measure of household indebtedness. I find that the elasticity of consumption with respect to income is approximately 21% higher for a one standard deviation increase in debt-to-income.

These overall elasticities fall towards the low end of the literature's estimates of elasticities of consumption in the face of persistent income shocks. These numbers often fall between 0.6 and 1 (elasticities for transitory income shocks are generally found to be approximately 0.2). In contrast, I find average numbers around 0.4 for what I find to be fairly permanent income shocks. One reason for the discrepancy is that I measure changes in spending at a quarterly frequency, while many previous studies use annual changes. Looking at year-on-year changes in spending and income, I obtain estimates about 50% larger, yielding overall elasticities of just over 0.6. Another reason for lower estimates is that the income shocks I measure, while persistent in at least the medium-term, may not be fully permanent, leading to a smaller consumption adjustment by households.

Finally, columns (7) and (8) give results demonstrating the applicability of the standard Bewley model to this data. I compare the direct impact of income changes on changes in consumption as well as the change in income interacted with household net wealth. I find similar effects in my data as in a simulated sample of households with income persistence calibrated to match my quarterly household financial data. I find evidence that increases in net assets drive down the sensitivity of consumption to income changes in a model with a persistent AR1 income process and a single savings bond mechanism.

Additional evidence for the mechanisms behind debt-induced household spending responses can be seen from the response of varying types of spending. Appendix Table B5 displays results examining several components of household spending: overall spending, spending on durables, and debt-related spending such as on credit card interest servicing or mortgage payments.[24] The specification follows those in Table 4, with both panel OLS and panel IV results in columns (1) - (3) and (4) - (6), respectively. Here, columns (1) and (4) report baseline results taken from Table 4. Consistent with the standard findings in the literature, I find that spending on durables is more elastic than spending as a whole and is also more sensitive to increased levels of debt.

Looking at debt-related spending, I find the opposite result. Debt-related spending is less elastic overall, as might be expected from fairly fixed payment terms, and is also much less related to household debt levels. That is, households with high levels of debt do not change debt-related spending relative to low-debt households following an income shock. This suggests that highly-indebted households cut non-

[24]Debt spending is measured as mortgage payments, loan payments, and credit card payments net of new credit card spending (eg. net credit card balance reductions).

debt spending to a much higher degree that do low-debt households, either prioritizing debt payments or a more limited ability to adjust debt payments on a high-frequency basis.

5.2 Robustness

I also employ alternate instruments for household leverage and other control variables in an effort to better isolate the causal impacts of household debt and leverage. First, I consider two alternate instrumentation strategies. The first dispenses homeowner interaction with the Saiz housing supply elasticity measure in the first stage of the instrument, somewhat decreasing first stage power. Column (1) of Table 5 gives the second stage results for this IV specification.

Column (2) similarly reports second stage results of a regression with a modified first stage. In this column, I allow for 'mechanically' time-varying levels of debt-to-asset ratios by including MSA time trends in the first stage of the IV regression. Here, I use the following first stage, wherein I regress time-varying debt-to-asset ratios on the Saiz housing supply elasticity measure as well as this measure interacted with MSA-specific quadratic time trends from 2008 to 2013:[25]

$$Leverage_{it} = \beta_0 + \beta_1 Saiz_i * MSATrend_{it} + \beta_2 Saiz_i * MSATrendSq_{it} + \beta_i HH_i + \beta_t Time_t + \xi_{it}$$

Column (3) reports results from a third IV specification. I utilize the Interstate Banking Restrictiveness Index which was developed by Johnson and Rice (2008), measuring the extent to which individual states restricted the ability of out-of-state banks to open in-state branches. This metric was used in Rice and Strahan (2010), among others, as an instrument for banking competition, finding that it drove down loan rates. I find that this instrument predicts higher levels of debt buildup among households in states with higher levels of banking competition. Using this instrument, in conjunction with the Saiz (2010) measure of housing supply elasticity, I find qualitatively similar results to those which instrument only for income and not assets or debt ratios.

In columns (4) and (5), I test two specifications that include other household level controls. Column (4) gives results from a version of the IV regression in Table 4, column (5), that also includes quadratic age terms as well as quadratic age terms interacted with changes in income. These terms help to control for variation in consumption elasticity that is driven by the age of the household. For example, younger households may have higher elasticities of consumption because a persistent change in household income represents a larger change in lifetime household assets. Finally, column (5) includes controls for household income as well as household income interacted with changes in income. In both of these columns, I find that, while the additional controls can explain some variation in household consumption elasticity, there remains a strong component of consumption variation related to household debt.[26]

[25] Also included in this first stage regression are all other applicable instruments including firm shocks and their interactions and were omitted for brevity.

[26] Results are robust to excluding all households residing in California, which exhibited the largest concentration of 'unconventional' mortgage loans in the nation. Similarly, utilizing the ratio of debt to liquid assets as a measure of indebtedness yields positive and highly significant results of a similar magnitude to the primary results.

5.3 Liquidity, Borrowing Constraints, and Debt

An important question is whether highly-indebted households cut back on consumption more strongly after income shocks because of liquidity or borrowing constraints.[27] Household models with precautionary motives to save generally feature increasing elasticities of consumption out of income as borrowing constraints get closer to binding. Failing to take household liquidity and borrowing constraints into account may erroneously attribute an increase in the sensitivity of household consumption to levels of leverage (debt-to-asset ratios).

For instance, households with little in the way of liquid assets may be more sensitive to income shocks and also have high levels of debt, in general. This may have come about through simply exhausting liquid savings and turning to debt because of unobserved shocks in the past or because of household attributes like age may jointly impact the amount of liquid assets accumulated and the amount of debt held by the household (eg. a young household has little liquid savings and also has newly acquired mortgage debt). In terms of credit access, a decline in asset prices, especially home prices, may have reduced households' ability to borrow against their home, thus forcing them to cut consumption. This would be most true for highly-leveraged households that had the least scope to further rely on home equity borrowing. Households with higher levels of debt may also have less access to additional consumer credit, even absent any decline in asset prices, in an environment where banks and credit card companies were increasingly wary about new lending. This lack of an ability to access consumer credit can decrease the ability of households to smooth consumption when subjected to an unanticipated negative income shocks.

In Table 6, I separately examine impacts of debt and assets, as well as effects of varying types of assets, on consumption responses. In all columns, I use firm shocks to instrument for changes in household income.[28] Column (1) displays results mirroring those from the baseline instrumental variables specification seen in Table 4. Columns (2) - (5) break down the debt to asset ratio to the household's balance sheet components.

Column (2) separately includes positive and negative balance sheet positions measured as the gross debt to income ratio and the gross asset to income ratio. I find that household spending is more sensitive to increases in debt than increases in assets, but that both balance sheet positions are significantly different than zero and are opposite signed. As the vast majority of households possess more assets than debt (only approximately 15% of households levels of debt exceed their assets), increases in both gross debt and gross assets push debt-to-asset ratios upwards towards 1. This points to one mechanism behind the increasing elasticity of consumption among households with higher levels of debt-to-assets. In addition, this result is consistent with the standard Bewley model, where increases in net assets reduce the elasticity of consumption.

Column (3) looks at household asset holdings in more detail, splitting assets into liquid and non-

[27]Possible alternate channels could include differential changes in beliefs about future income or labor market opportunities or about future credit availability, mental accounting/rule of thumb budgeting, or leverage and debt being directly in a households' utility function. For instance, if a household thinks it is wise to hold debt equal to less than 2 years of annual income, a decline in income may prompt a pay-down of debt even in the absence of binding borrowing constraints.

[28]I normalize debt and asset measures by household income in 2008 given that wealth's smoothing effects are intimately linked to current household income and spending.

liquid assets. Non-liquid assets include vehicles and home values, while liquid assets include bank balances, reported cash holdings, and equities. Here, I find similar, though opposite signed, coefficients on both the interaction between changes in income and liquid assets as well as the interaction of income changes with debt. I find a much smaller coefficient on the interaction term using non-liquid assets. These results indicate that a $1 increase in both illiquid assets and debt would yield a small increased sensitivity to household income changes. However, borrowing $1 and holding it as a liquid asset yields a combined effect in the opposite direction, with household consumption sensitivity dropping significantly.

Given that the years leading up to the Great Recession featured large increases in debt used to finance additional housing asset purchases, these results expose a primary driver of increasing aggregate consumption elasticity with respect to income. Even without any decline in asset prices, the increasing levels of household wealth being held in illiquid assets would be predicted to increase consumption elasticity. Kaplan and Violante (2011) outline a framework that could rationalize differential responses to gross debt and gross asset holdings in an incomplete markets model and which is consistent with these results. They introduce two asset classes, bonds and housing, where housing provides higher returns but is less liquid, requiring transaction costs to convert into consumption. They find that this framework can lead to the existence of "wealthy hand-to-mouth" households which possess substantial amounts of net assets but who still respond strongly to income shocks due to a lack of liquid assets available to smooth consumption.

Columns (4) and (5) utilize changes in logged non-durable and durables expenditures as dependent variables rather than all spending. I find generally weaker responsiveness of non-durables spending and higher responsiveness of durables spending.

In Table 7, I examine the impact of debt on consumption responses after controlling for a range of measures of credit access and liquidity. All additional variables are measured at a household level in 2008 and held constant throughout the sample period. Each continuous variable is normalized such that it is of standard deviation 1.

Column (2) includes individual credit scores, finding a negative relationship between higher credit scores and consumption responses. This finding is driven by the fact that high credit score households will tend to have access to larger amounts of additional and unobservable credit usable for consumption smoothing. Columns (3) and (4) include interactions with 'unused credit' and liquid savings. 'Unused credit' represents room to borrow on credit cards and from HELOC loans that is untapped.[29]Liquid assets, as in Table 6, represents money held in savings and checking accounts, cash, and non-retirement equity accounts. For both of these interactions I find negative effects, as liquid savings and readily available credit enable households to smooth consumption.

Column (5) includes an interaction of changes in income with one attempt to measure changes in credit supply. While I generally cannot observe the menu of options that a household faces when desiring additional credit, I can observe some changes in existing household credit limits. Increases in credit limits are generally demand driven, with households requesting additional credit, while decreases

29

[29]Unused credit is measured as the remaining room to borrow on existing credit cards in addition to open lines of credit that the households have access to. Credit cards with no pre-set borrowing limit are measured as having a $200,000 credit limit, equal to the 99.9th percentile of credit cards.

in limits may be indicators of shifts in credit supply. I interact changes in income with an indicator that the household had at least one of its credit limits cut in the previous year. This measure is noisy given the relative rarity of credit limit declines faced by the sample population. However, I find some evidence that having credit limits cut induces greater subsequent consumption responses to income shocks.

Finally, I also interact changes in income with the highest interest rate faced by households. As in Kreiner, Lassen, and Leth-Petersen (2014), the marginal interest rate on a household debt account may be indicative of the potential credit supply that the household faces. More constrained and less credit-worthy borrowers would possess accounts with higher interest rates as they exhaust cheaper sources of borrowing. Again, I find a marginally significant value here, with households possessing higher interest rates on their debt being more sensitive to changes to their income.

Figure 9 gives a clear indication of the importance of liquidity constraints to this relationship. This graph shows results from a nonparametric version of the regression in column (6) of Table 7, with the debt to asset ratio being split into 10 variables by debt decile. The baseline regression excludes all credit and liquidity covariates, finding a strongly positive relationship between the debt to asset ratio and the responsiveness of consumption to a given change in income. The second line includes all previously mentioned credit and liquidity covariates. Here I find a near-zero relationship, implying that households possessing similar access to liquid assets and credit behave similarly when subjected to a change in income, even if these households possess different levels of debt or leverage.

While the measures of credit and liquid savings used here are imperfect, as they cannot fully control for all aspects of the potential access to credit, I argue that they provide a good measure of immediate access to credit and liquid savings. For instance, households with high credit scores maintained greater access to credit markets even during the Great Recession and households with ample liquid savings could self-finance spending. In addition, I find that high credit scores are correlated with the presence of more liquid wealth and observable borrowing capabilities, suggesting that to the extent that measurement error is present, it is largely understating the degree to which high credit households are unconstrained in their borrowing abilities. Moreover, the time period I study was largely one of reduced credit access for households, meaning that there was less of an ability to obtain new lines of credit, credit cards, or HELOCs, even if a household desired to do so.

Another test of the impact of liquidity constraints is to examine differential impacts of debt across different types of income shocks. Table 8 displays results when solely looking at positive or negative shocks to household income. Columns (2)-(3) and columns (4)-(5) look at different consumption responses across heterogeneous households, in terms of either debt to income ratios or debt-to-asset ratios, to solely negative or solely positive shocks. I find much smaller consumption responses to positive income shocks than to negative shocks. However, I find that even following positive shocks to household income, more highly-indebted households exhibit higher elasticities of consumption with respect to income than do low debt households (albeit with marginal significance). This finding is inconsistent with a purely liquidity or borrowing constraints story, as households always have the ability to smooth consumption through saving in the presence of unexpected positive innovations to their income process.

Finally, I also note that results differ across types of debt when examining spending elasticities' relationship to household debt. Most notably, I find that there is a stronger relationship between spending elasticity and revolving credit card debt than with mortgage debt (both gross and net mortgage debt). Again, this divergence appears to be largely driven by liquidity and credit. That is, households who possess large amounts of mortgage debt often still have access to substantial liquid savings that allow them to smooth consumption. That presence of significant credit card debt generally indicates that most liquid savings have been exhausted (else households would usually find it advantageous to pay off the high-interest credit card debt) and cannot easily smooth unexpected changes in consumption.

Overall, these results indicate that the primary reasons consumption responses are higher among highly indebted households are borrowing and liquidity constraints. That is, household debt and leverage generally did not drive consumption behavior in and of themselves, but households that were highly leveraged were also very constrained in terms of their liquid savings and available credit. However, the presence of statistically significant effects of higher levels of debt in some specifications, among even the most liquid households and following positive changes in income, suggests that some of this relationship may be due to alternate channels. Such channels could include differential changes in expectations about future income or behavioral explanations in which households directly target levels of debt as a function of household income.

6 Aggregate Effect Estimates

One important statistic that can be derived from the results of this paper is an estimate of the aggregate effect of household debt during the Great Recession. Given that increased household debt drives a steeper decline in consumption in response to decreases in household income, it is useful to note to what extent the higher levels of debt among households immediately prior to the Great Recession drove a larger drop in consumption than would otherwise have been seen.

From column (3) from Table 6, combined with measures of the distribution of household balance sheet positions, I derive a gauge of the impact of heightened debt and illiquid assets, primarily in the form of home equity and mortgage debt, on consumption. Figure 6 displays the distribution of household debt-to-income ratios in 1983 and 2007, as measured by the Survey of Consumer Finance.

Assuming that the entirety of the observed effects of debt-to-income ratios is causal, I utilize data on a representative sample of households from the SCF to construct household debt, illiquid assets, and liquid asset distributions for both the 1983 and 2007 populations. I then subject each population distribution to the observed average decline in household income seen during the Great Recession, which is approximately 8%.[30]

With these parameters, I estimate a 3.29% decline in consumption under the assumption that the population has the joint debt and asset distribution seen in 1983, and a 4.16% decline in consumption with the actual debt distribution seen in 2007.[31] Thus the estimates from this analysis suggest a potentially sizable impact of household debt in triggering larger declines in consumption, with the aggregate effect on the order of a 25% larger fall in consumption. This consumption decline only incorporates the decline in household income during the Great Recession; the decrease in consumption would be

[30] Household income decline from data from the US Dept of Commerce, Census Bureau, using the top to the bottom of the market: 2007 to 2011.

substantially larger when incorporating the concurrent asset price declines during this period.

Similarly, I can utilize the point estimates obtained from an estimation considering only negative household income shocks (similar to the exercise undertaken in Column (3) of Table 8, but including both liquid and illiquid asset ratios), given that the income shock suffered by American households during the Great Recession was a negative one. Using these estimates, I find that the 1983 households would see an average decline in consumption of about 4.55% while 2007 households see a decline of 6.10%, a 34% larger decrease.

7 Conclusion

The Great Recession of 2007-2009 featured a large increase in household debt in the years leading up to its onset and a large decline in household spending during the recession. These two features have led to a growing number of economists and policymakers linking these two features, describing a recession in which highly-indebted households were hit by financial shocks, such as asset price or income declines, and subsequently cut back on consumption while attempting to repair damaged balance sheets.

I evaluate this claim, examining the role that debt played in driving the decline in consumption seen during the Great Recession and providing new empirical evidence regarding the relationship between household balance sheets, income, and spending. In doing so, this paper provides a fuller characterization of the interaction between household balance sheets and other household financial decision-making which is crucial to understanding the microeconomic underpinnings of business cycles.

To this end, I employ a dataset with comprehensive financial information on over 150,000 American households from 2008 to 2013. This data covers household debt and assets as well as information regarding income and consumption. In addition, I link households to their employers. I utilize variation in both asset and debt holdings as well as innovations in the household income process in order to test whether balance sheet positions drove heterogeneous consumption responses to a given change in income. To address endogeneity concerns, I use shocks originating from employers as well as drivers of balance sheet positions such as geographically determined housing supply elasticities and inter-state banking regulations as instruments for household income and balance sheet positions.

This paper presents three primary results. Using a new source of matched employer-employee data that comprehensively tracks the entirety of household spending, assets, income, and debt, I find that households are significantly impacted by shocks to their employers such as large and surprising earnings reports, layoff announcements, and large write-offs by the firm. I find that these shocks are unanticipated by households and exhibit a high degree of persistence.

Consistent with standard incomplete-markets models, I find that larger net asset holdings mute consumption elasticities. In a departure from these models, I also find that the elasticity of consumption with respect to income among highly-indebted households is significantly higher than in low debt households even after controlling for net assets. In conjunction with the distribution of debt immediately preceding the recent recession as well as the distribution in earlier years, I estimate that higher

[31] Similar estimates utilizing the leverage distribution in 2001 yield an aggregate fall in consumption of approximately 3.75%.

levels of debt significantly worsened the decline in household consumption across the country.

Finally, I show that credit constraints play a dominant role in driving differential household consumption responses across households with varying levels of debt. When controlling for various measures of access to liquid assets and consumer credit, the coefficients on the interaction between changes in income and household leverage generally decline into marginal significance or insignificance. In addition, I find much lower effects of debt and leverage following positive shocks to income, consistent with a credit constraints explanation. Consistent with prior work, I find a similar pattern holds when examining asset price changes.

However, in some specifications, household debt remains a significant predictor of household consumption behavior, even after conditioning on assets and credit. This result implies that the causal relationship between household leverage and consumption responses may not be entirely an artifact of increased borrowing and liquidity constraints. One potential alternate channel is that households have an aversion to holding more debt than a household-specific target, thus causing more highly-indebted households to adjust consumption to a greater degree following income declines in order to maintain a target ratio of debt to income or assets.

In total, this paper's findings suggest that changes in household balance sheets, in particular increases in levels of household debt, have been important drivers of household behavior during the Great Recession and subsequent recession. Much of this effect was driven by tightened borrowing constraints and increases in the level of non-liquid assets relative to liquid assets. The buildup in household debt in the years leading up to the onset of the recession increased sensitivity of household spending to both income and asset shocks, causing heterogeneous responses to income and asset price declines. Moreover, this paper points to the importance of models which incorporate heterogeneous households and more detailed accounts of household balance sheet positions.

References

[1] Agarwal, S., C. Liu, and N. S. Souleles. *The reaction of consumer spending and debt to tax rebates.* Journal of Political Economy, 2007.

[2] Aguiar, Mark and Erik Hearst. *Consumption versus Expenditure.* Journal of Political Economy, 2005.

[3] Alan, Sule, Thomas Crossley, and Hamish Low *Saving on a Rainy Day, Borrowing for a Rainy Day.* IFS Working Paper W12/11, 2012.

[4] Bishop, Thomas and Cheolbeom Park. *Borrowing Constraints, the Marginal Propensity to Consume, and the Effectiveness of Fiscal Policy.* Korean Development Institute (KDI) Journal of Economic Policy, 2011.

[5] Blundell R, Pistaferri L, Preston I. *Consumption inequality and partial insurance.* American Economic Review, 2008.

[6] Carroll, Christopher D. and Wendy Dunn. *Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets.* NBER Macroeconomics Annual, 1997.

[7] Chmelar, Ales. *Household Debt in Europes Periphery: The dangers of a prolonged recession.* ECRI Commentary No. 12, 2012.

[8] Davis, Steven J. and Til von Wachter. *Recessions and the Costs of Job Loss.* Brookings Papers, 2011.

[9] Dynan, Karen *Is a Household Debt Overhang Holding Back Consumption?.* Brookings Papers on Economic Activity, 2012.

[10] Eggertsson, Gauti B. and Paul Krugman *Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo approach.* Quarterly Journal of Economics, 2012.

[11] Flavin, Marjorie A. *The Adjustment of Consumption to Changing Expectations About Future Income.* Journal of Political Economy, 1981.

[12] French, Eric, Taylor Kelley and An Qi *Expected income growth and the Great Recession.* Economic Perspectives, 2013.

[13] Glaeser, Edward, Joseph Gyourko and Albert Saiz *Housing Supply and Housing Bubbles.* NBER Working Paper No. 14193, 2008.

[14] Glick, Reuven and Kevin J. Lansing. *Global Household Leverage, House Prices, and Consumption.* San Francisco Fed mimeo, 2010.

[15] Gropp, Reint, John Krainer, and Elizabeth Laderman. *Did Consumers Want Less Debt? Consumer Credit Demand Versus Supply in the Wake of the 2008-2009 Financial Crisis.* San Francisco Fed working paper, 2010.

[16] Guerrieri, Veronica and Guido Lorenzoni *Credit Crises, Precautionary Savings, and the Liquidity Trap.* NBER Working Paper, 2011.

[17] Hall, Robert and FS Mishkin *The sensitivity of consumption to transitory income: estimates from panel data on households.* Econometrica, 1982.

[18] Hall, Robert E. *The Long Slump.* American Economic Review, 2011.

[19] Igan, Deniz, Daniel Leigh, John Simon, and Petia Topalova. *Dealing with Household Debt.* IMF mimeo, 2012.

[20] Isaksen, Jacob, Paul Lassenius Kramp, Louise Funch Srensen, and Sren Vester Srensen. *Household Balance Sheets and DebtAn International Country Study.* Denmarks Nationalbank, Monetary Review, 2011.

[21] Jacobson, Louis S., Robert J LaLonde and Daniel G. Sullivan *Earnings Losses of Displaced Workers.* American Economic Review, 1993.

[22] Jappelli, Tullio and Luigi Pistaferri *The Consumption Response to Income Changes.* Journal of Political Economy, 2005.

[23] Johnson, Christian A. and Tara Rice *Assessing a Decade of Interstate Bank Branching.* Working Paper Series WP-07-03., 2008.

[24] Kaplan, Greg and Giovanni Violante. *A Model of the Consumption Response to Fiscal Stimulus Payments.* NBER Working Paper Series, 2011.

[25] King, Mervyn *Debt deflation: Theory and evidence.* European Economic Review, 1994.

[26] Koo, Richard *Balance Sheet Recession: Japan's Struggle with Uncharted Economics and its Global Implications.* John Wiley and Sons, 2003.

[27] Kreiner, Claus T. and David Dreyer Lassen and Sren Leth-Petersen *Liquidity Constraint Tightness and Consumer Responses to Fiscal Stimulus Policy.* University of Copenhagen Working Paper, 2014.

[28] Krueger, Dirk and Fabrizio Perri *Understanding Consumption Smoothing: Evidence from the US Consumer Expenditure Survey.* Journal of the European Economic Association Papers and Proceedings, 2005.

[29] Krueger, Dirk and Fabrizio Perri *How do households respond to income shocks?.* University of Pennsylvania mimeo, 2009.

[30] Lazear, Edward and Paul Oyer *Personnel Economics.* Handbook of Organizational Economics Chapter, 2010.

[31] Leamer, Edward *Housing is the business cycle.* NBER Working Paper No. 13428, 2007.

[32] Mian, Atif R., and Amir Sufi *Household Leverage and the Recession of 2007 to 2009.* IMF Economic Review. https://www.imf.org/external/np/res/seminars/2009/arc/pdf/mian.pdf, 2010.

[33] Mian, Atif R. and Amir Sufi *House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis.* American Economic Review, 2011.

[34] Mian, Atif R., Kamelesh Rao and Amir Sufi *Household Balance Sheets, Consumption, and the Economic Slump.* Quarterly Journal of Economics, 2013.

[35] Mian, Atif R., and Amir Sufi *What Explains High Unemployment? The Aggregate Demand Channel.* Econometrica, 2013.

[36] Midrigan, Virgiliu and Thomas Philippon *Household Leverage and the Recession.* NBER Working Paper No. 16965, 2012.

[37] Mishkin, Frederic *The Household Balance Sheet and the Great Depression*. Journal of Economic History, 1978.

[38] Mishkin, Frederic *What Depressed the Consumer? The Household Balance Sheet and the 1973-75 Recession*. Brookings Mimeo, 1977.

[39] Olney, Martha L. *Avoiding Default: The Role of Credit in the Consumption Collapse of 1930*. Quarterly Journal of Economics, 1999.

[40] Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. *The short-and long-term career effects of graduating in a recession*. American Economic Journal: Applied Economics, 2012.

[41] Petev, Ivaylo, Luigi Pistaferri and Itay Saporta Eksten *An Analysis of Trends, Perceptions, and Distributional Effects*. The Great Recession, Russell Sage Foundation, New York, pp. 161-95., 2011.

[42] Pistaferri, Luigi and Itay Saporta-Eksten. *Changes in the Income Distribution and Aggregate Consumption*. Stanford Working Paper, 2012.

[43] Rice, Tara and Philip E. Strahan. *Does Credit Competition Affect Small-Firm Finance?* Journal of Finance, vol. 65, no. 3, pp. 861-889, 2010.

[44] Roys, Nicolas *Persistence of Shocks, Employment and Wages*. University of Wisconsin-Madison mimeo, 2013.

[45] Saiz, Albert. *The Geographic Determinants of Housing Supply*. Quarterly Journal of Economics, 2011.

[46] Shea, J. *The reaction of household consumption to predictable changes in social security taxes*. American Economic Review, 1995.

[47] Souleles, Nicholas S. *The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans*. American Economic Review, 1999.

[48] Wilcox DW *Social security benefits, consumption expenditure, and the life cycle hypothesis*. Journal of Political Economy, 1989.

[49] Wolpin, KI *A new test of the permanent income hypothesis: the impact of weather on the income and consumption of farm households in India*. International Economics Review, 1982.

[50] Zeldes, SP *Consumption and liquidity constraints: an empirical investigation*. Journal of Political Economy, 1989.

[51] Zillow.com *Average Home Price Data by Zip Code-Year*. Zillow.com Database, 2012.

Figure 1: **Ratio of Household Debt to Disposable Income and Household Debt to Assets**

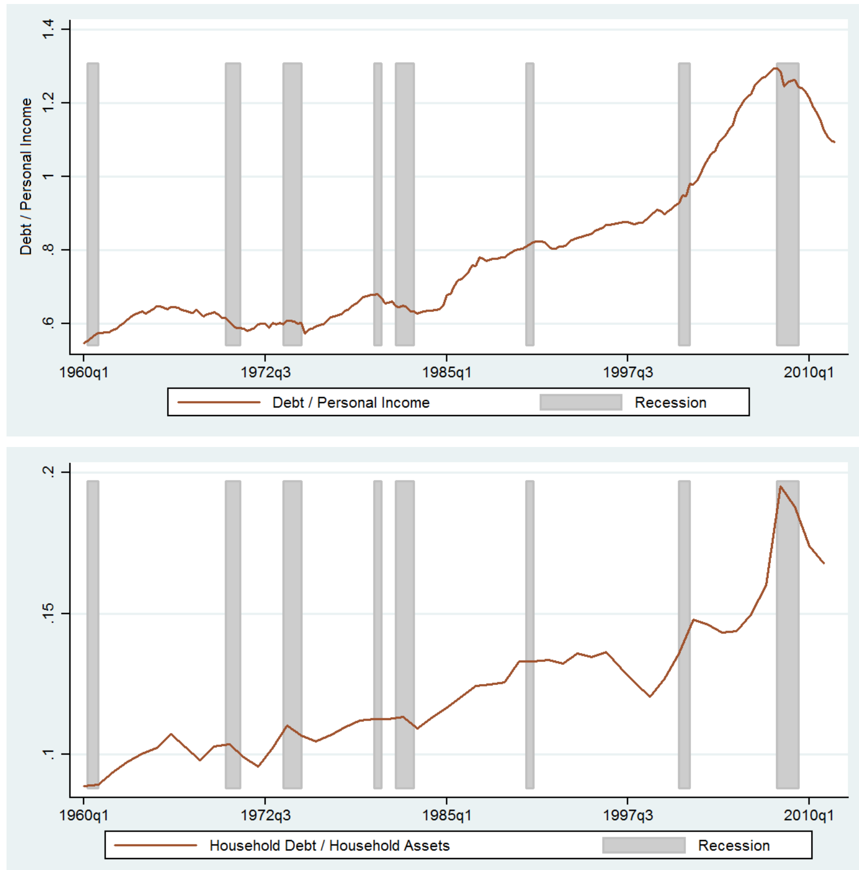

Notes: Indexes of household debt, household assets, and household disposable income are taken from the Federal Reserve Economic Data. Disposable income is series DSPI which is total nominal disposable household income, seasonally adjusted at an annual rate. Debt represents series CMDEBT, total nominal seasonally adjusted. household and nonprofit credit market instrument liabilities. Assets represents series HNOTASA, including total household and nonprofit asset levels, at an annual level.

Figure 2: **Firm Shocks Over Time**

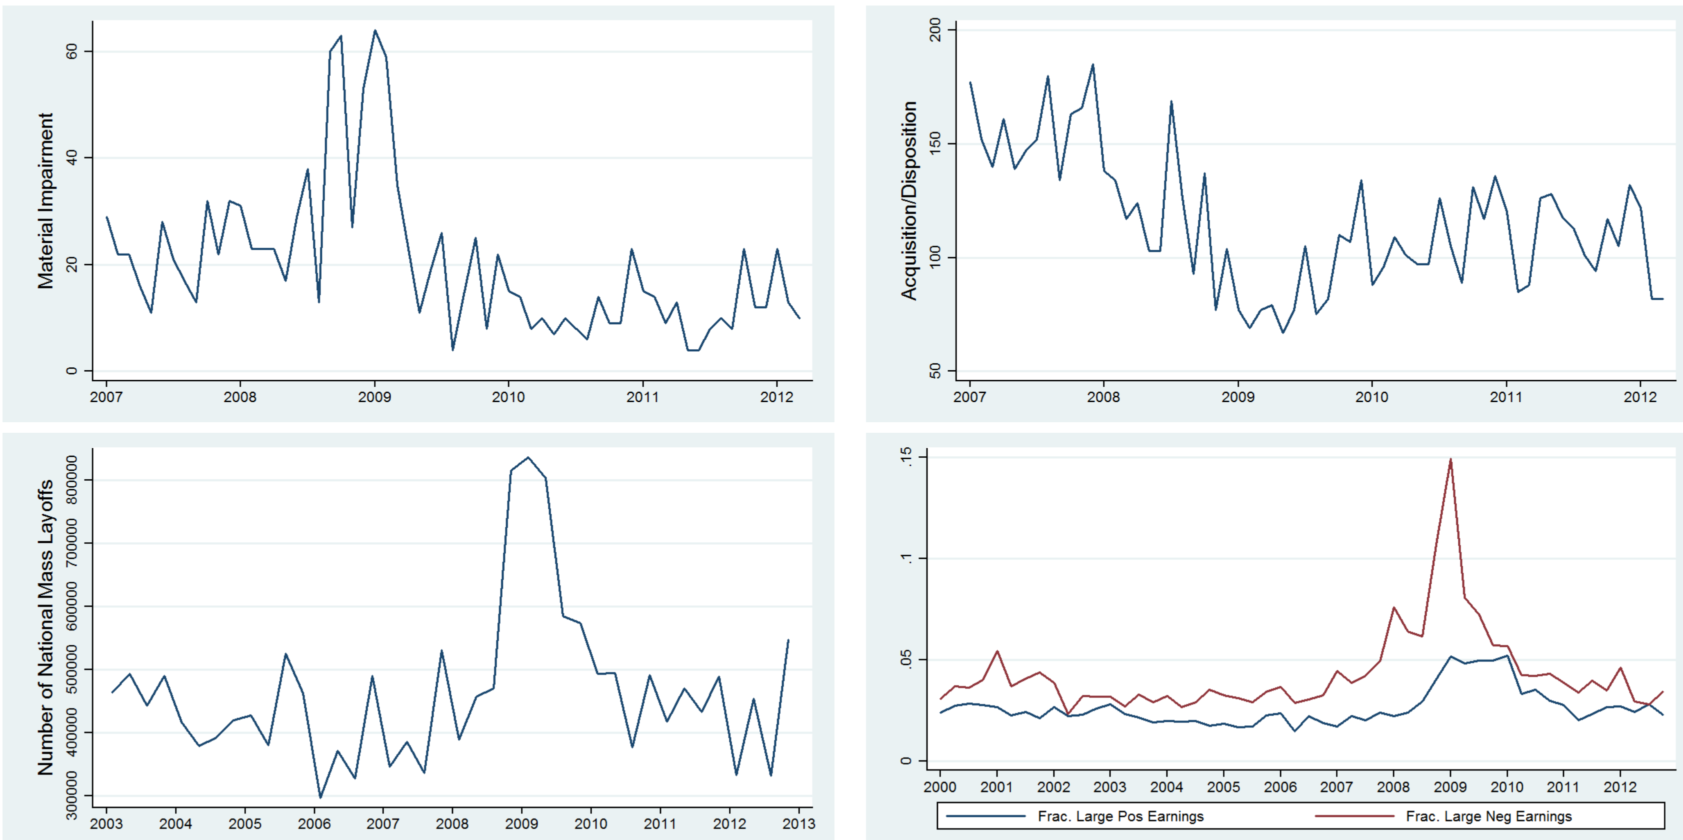

Notes: Clockwise from top-left: (1) Total material impairments by quarter from the first quarter of 2007, from SEC 8K data. (2) Acquisition and Disposition Completion data from the first quarter of 2007, from SEC 8K data. (3) Total number of layoffs by quarter, from Mass Layoff series from the BLS. (4) 'Large' surprising earnings per share reports, from IBES reports where an earnings report is considered a large surprise if the predicted earnings per share are more than 2% of the share price from the true earnings per share. Positive and negative surprises are tallied separately.

Figure 3: **Change in Log(Income) Before and After Firm Shocks**

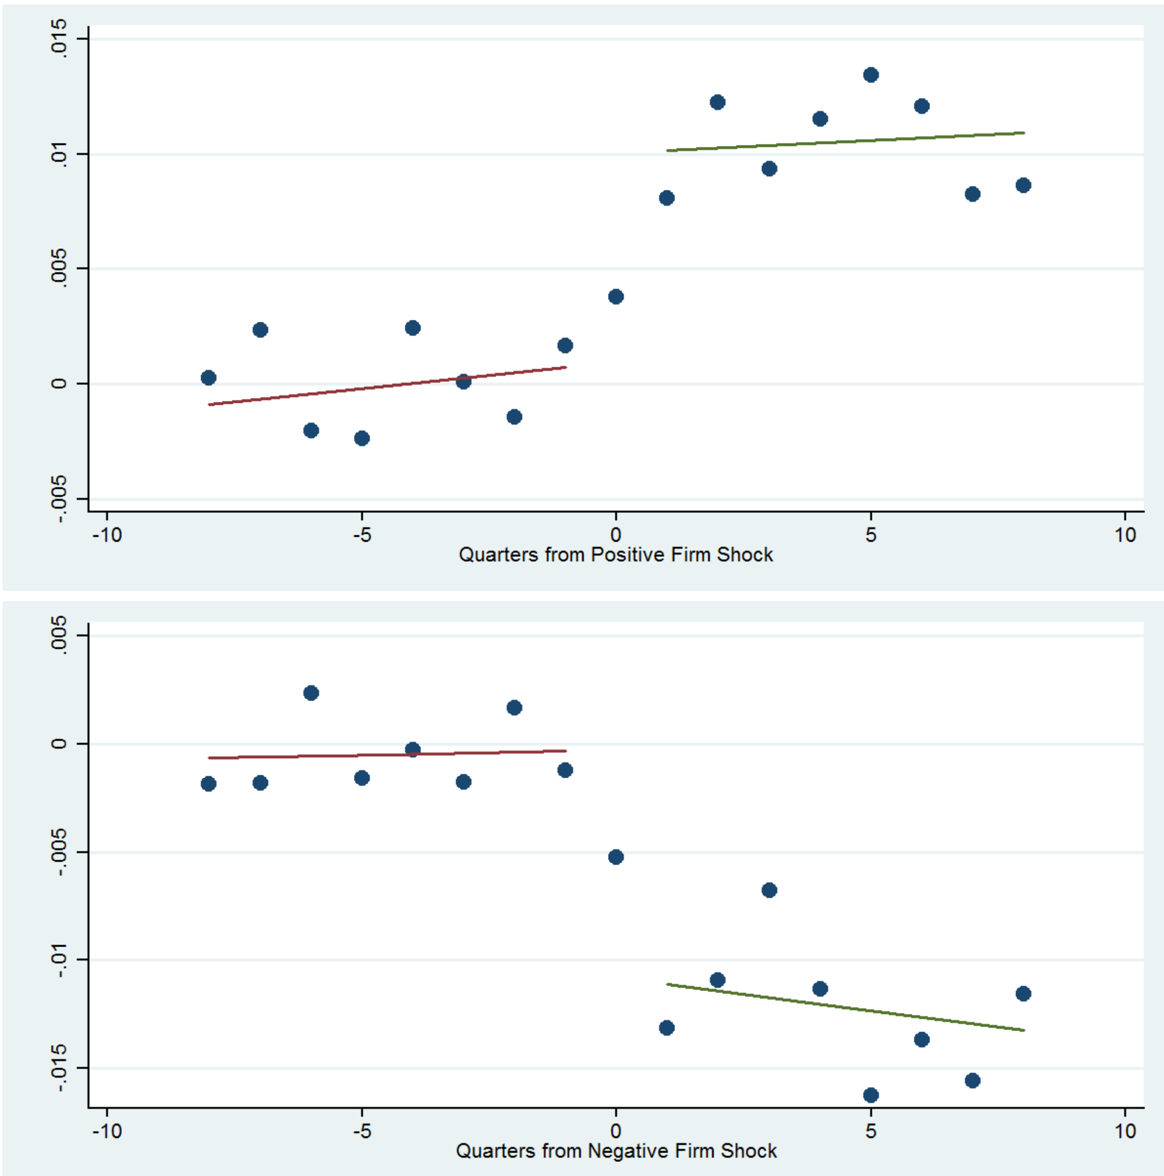

Notes: Data derived from logged average quarterly income across all households experiencing firm shocks. Data extends 6 quarters prior to each shock and 6 quarters following each shock. A time trend is removed from the data based on the pre-shock trend (across 6 quarters prior to shock) and residuals are plotted.

Figure 4: **Wide Variation in Household Income (Blue Fill) and Assets Among Employees of Each Firm (Red Empty)**

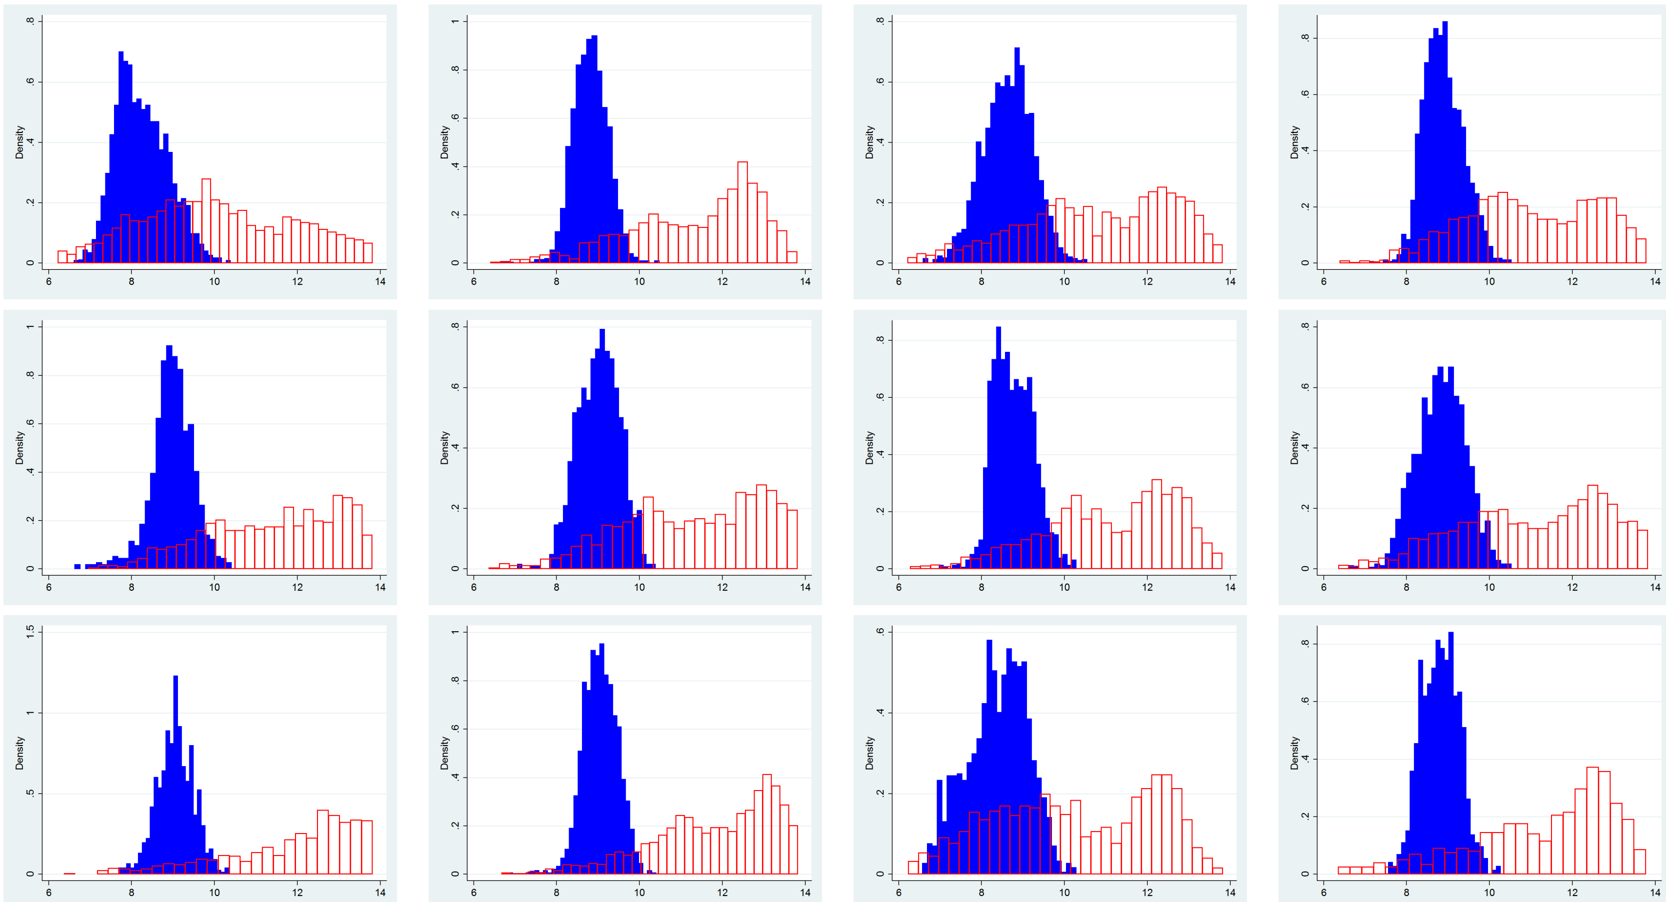

Notes: Figures display dispersion in logged income (income from all sources including paychecks, rental income, dividends, etc.) and logged total financial assets (bank balances, cash, home value, equity balances, etc.) across individual households for a given firm. Both total financial assets and average monthly income are censored at the 99.5th percentile for display purposes. The 12 firms with the highest number of users are displayed, with the number of users ranging from 6355 to 16464.

Figure 5: **Comparison Between User-Derived Spending and Census Retail Data, by Category**

Notes: From top-left, by row: (1) Per-Household monthly spending on gasoline from online financial data and Census Retail per-household spending on gasoline. (2) Per-Household monthly spending on restaurants and food away from home from online financial data and Census Retail per-household spending on restaurants. (3) Per-Household monthly spending on furniture from online financial data and Census Retail per-household spending on furniture. (4) Per-Household monthly spending on electronics from online financial data and Census Retail per-household spending on electronics. (5) Per-Household monthly spending on groceries from online financial data and Census Retail per-household spending on groceries. (6) Per-Household monthly spending on clothing from online financial data and Census Retail per-household spending on clothing. (7) Per-Household monthly spending on building supplies and hardware from online financial data and Census Retail per-household spending on building and home supplies. (8) Overall per-household monthly spending on motor vehicles from online financial data and overall Census Retail per-household spending on motor vehicles. (9) Overall per-household monthly spending on arts, books, and hobby supplies from online financial data and overall Census Retail per-household spending on hobby supplies. Each graph's y-axis represents household spending per month scaled to average 100 over the sample period.

Figure 6: **Ratio of Total Debt to Total Household Income**

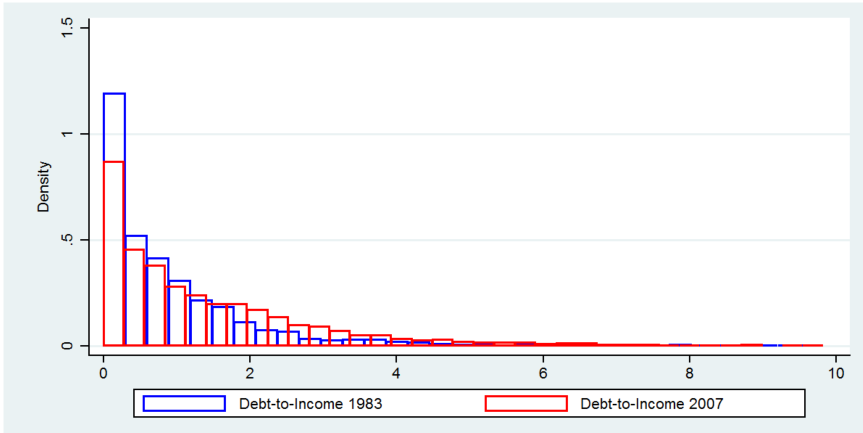

Notes: Data from the 1983 and 2007 Surveys of Consumer Finances. Each observation represents a household with a given total debt to gross income ratio. Total Household Debt and Debt-to-income ratios are censored at 10 for display purposes.

Figure 7: **Average Observed Household and Zillow.com Home Prices**

Notes: 'Zillow' data from Zillow.com, tracking average zip-code level home prices by month. Zillow data aggregated to yearly level by zip-code, then logged. User data taken by averaging real estate values by year within each self-identified homeowner and then averaging across all homeowners by zip-code, then taking logs. Line of best fit shown. Correlation is 0.881.

Figure 8: **Bank Savings and Credit Distributions, Observed and Survey of Consumer Finance (2010)**

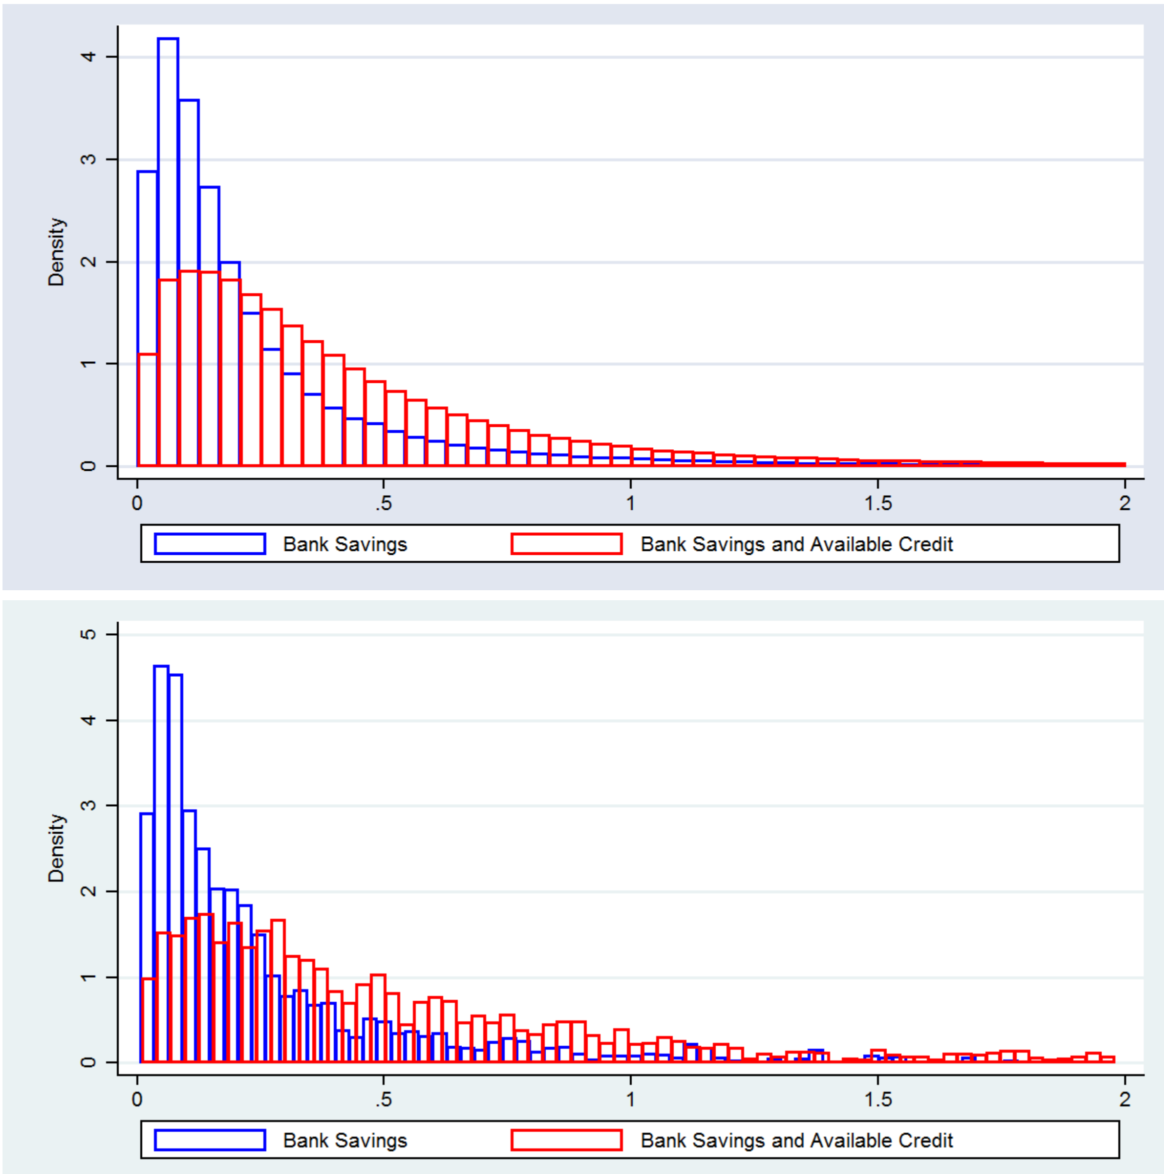

Notes: Graphs represent distributions of bank savings and available consumer credit among all households in personal financial website sample and among all households in the 2010 SCF. In both cases, bank savings include all bank holdings such as balances of savings accounts, checking accounts, money market accounts, etc. as well as the available remaining balance on credit cards. Ratios of these assets to annual household income are taken. Both distributions censored at 2 for display purposes.

Figure 9: **Consumption Elasticity with Respect to Income Across Debt/Asset Deciles**

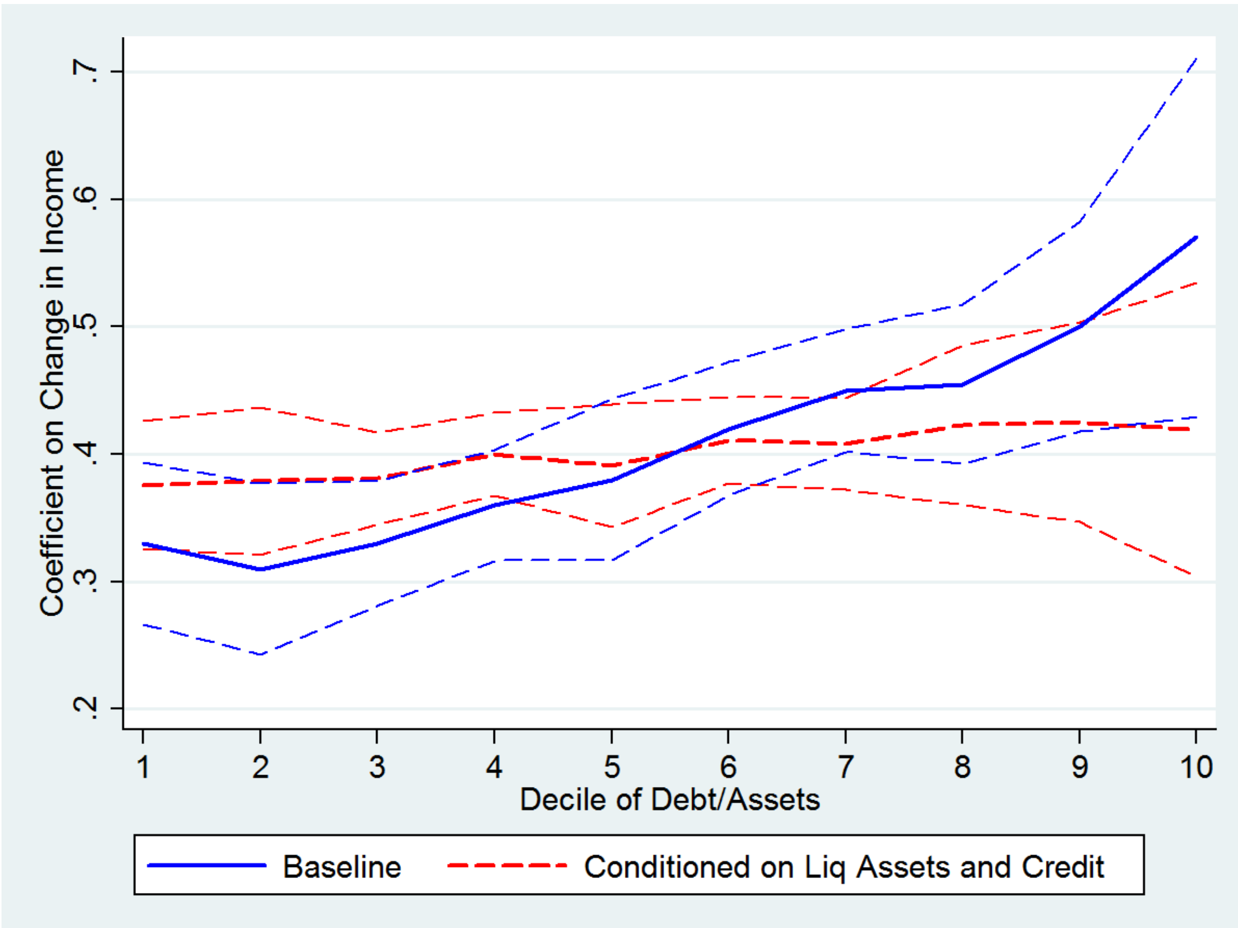

Notes: Debt/Asset calculated as total household debt divided by total household assets. Households divided into deciles by 2008 value of debt/assets. Baseline coefficients represent coefficients on panel fixed effects regression of the change in log consumption on the change in log income, the change in log income interacted with decile of debt/assets, along with household and time fixed effects. Conditioning on Liquid Assets and Credit Scores coefficients runs the same regression including interactions of changes in income with liquid assets and the credit score and then assigning all households the average credit score and level of liquid assets for scaling.

Table 1: **Summary Statistics (CPS Weighted)**

	Median	Mean	Std Dev	Number
Income (Monthly)	\$4,810	\$6,758	\$6,287	-
Spending (Monthly)	\$4,790	\$6,725	\$6,156	-
Equity Wealth	\$902	\$42,790	\$110,158	-
Property Value (Homeowners)	\$232,500	\$316,801	\$275,256	-
Property Value (Renters)	\$11,000	\$54,897	\$152,626	-
Total Net Wealth	\$20,990	\$127,414	\$280,266	-
Total Debt to Assets	0.34	0.68	1.14	-
Housing Debt to Assets	0.46	0.48	0.438	-
Gross Debt to Income	0.24	0.83	1.31	-
Num Households	-	-	-	156,604
Num Matched Firms	-	-	-	361
Num Large Earnings Rep.	-	-	-	82
Num Layoffs	-	-	-	79
Num Completed Acq/Disp	-	-	-	39
Num Material Impairments	-	-	-	37

Notes: Data for the 156,604 users who form the sample group. Each user meets a variety of characteristics such as having been an active site user for more than 12 months, being able to be matched to an employer, having more than a single account linked, etc. Households are weighted on four characteristics (age, sex, income, and state of residence) in order to match the observed CPS distribution of heads of households in the United States. Positive and Negative Earnings Surprises are indicators whether an earnings report differed from the mean expectation of earnings (from IBES data) by more than 2% of the stock's value. Layoffs, Completed Acq/Disp, and Material Impairments are indicators derived from SEC 8K data. 'Positive' shocks are Positive Earnings Surprises and Completed Acq/Disp while the remaining firm shocks are classified as 'Negative' shocks. Time period covered is January 2008 - December 2013.

Table 2: **Tests for Anticipation of Firm Shocks**

	(1) Firm Returns	(2) Layoffs	(3) Pos Earnings	(4) Neg Earnings	(5) Lagged Firm Returns	(6) Forward Firm Returns
Lagged Income	-0.084	-0.045	0.025	-0.063		
	(0.220)	(0.126)	(0.0712)	(0.103)		
Lagged Wealth	-0.0591*	0.0169	0.0161	0.0241		
	(0.0307)	(0.0131)	(0.0140)	(0.091)		
Lagged Spending	-0.0447	0.0393	-0.0386	0.0239		
	(0.0601)	(0.0591)	(0.0602)	(0.0629)		
Postive Earnings Surprise					-0.0201	0.0647**
					(0.0152)	(0.0250)
Negative Earnings Surprise					-0.0165	-0.0422***
					(0.0159)	(0.0143)
Layoffs					-0.0118	-0.0142*
					(0.0105)	(0.00789)
Observations	7,786,804	7,786,804	2,157,975	2,157,975	7,560	7,033
Level of Analysis	HH-Month	HH-Month	HH-Month	HH-Month	Firm-Month	Firm-Month
R^2	0.429	0.085	0.372	0.224	0.391	0.393
Period and Firm FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	NO	NO

*** p<0.01, ** p<0.05, * p<0.1

Notes: All lagged (forward) variables refer to the three months prior (subsequent) to a given date. Income, wealth, and spending denote logged total income, logged total wealth (property, bank balances, etc.), and logged spending. Firm returns are cumulative monthly returns. Layoffs and earnings surprises are binary indicators for the month where the event occurred. Earnings report data taken from IBES reports. Layoff announcements are taken from SEC 8K layoff and closure notifications. Columns (3) and (4) run on only months in which there was an earnings report issued. Columns (5) and (6) are run at a firm-month level. Regressions span January 2008-December 2013. All standard errors clustered at a household level.

Table 3: **Firm Shocks Drive Employee Income and Bonuses**

	(1)	(2)	(3)	(4)	(5)	(6)
	ΔLog(Inc)	ΔLog(Inc)	ΔLog(NonFirm)	ΔLog(Paycheck)	ΔScaled Bonus	Log(Income)
Postive Earnings Surprise	0.0121***	0.0099***	0.00171	0.00575***	0.234***	
	(0.0041)	(0.0024)	(0.0117)	(0.0016)	(0.0611)	
Negative Earnings Surprise	-0.0196***	-0.0134***	-0.0037	-0.0118***	-0.308**	
	(0.0066)	(0.0041)	(0.0120)	(0.0038)	(0.0734)	
Log(Number Layoffs)	-0.00178***	-0.00201***	-0.00013	-0.00144***	-0.0311***	
	(0.00038)	(0.00041)	(0.0025)	(0.00031)	(0.0038)	
Acquisition/Disposition	0.0145***	0.0079**	0.0016	0.0065***	0.291**	
	(0.00704)	(0.00234)	(0.00351)	(0.00218)	(0.0813)	
Material Impairment	-0.0131***	-0.0128***	-0.0034	-0.0064***	-0.246***	
	(0.00556)	(0.00388)	(0.0085)	(0.0024)	(0.0781)	
Positive Shock						0.01114***
						(0.00195)
Negative Shock						-0.01389***
						(0.00240)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
R^2	0.263	0.283	0.249	0.319	0.302	0.243
Period FE	YES	YES	YES	YES	YES	YES
Household FE	NO	YES	YES	YES	YES	YES

*** p<0.01, ** p<0.05, * p<0.1

Notes: Positive and Negative Earnings Surprises are indicators whether an earnings report differed from the mean expectation of earnings (from IBES data) by more than 2% of the stock price. Acq/Dis of Assets and Material Impairments are indicators derived from SEC 8K data. Layoffs are the number of layoffs derived from scraping AccessWorldNews newspaper database. NonFirm income is all non-paycheck and non-bonus income. Column (6) condenses the various firm shocks into 'Positive' and 'Negative' binary indicators, with Positive Earnings Surprises and Completed Acq/Disp classified as positive shocks and the rest as negative shocks. Regressions span January 2008-December 2013. All standard errors clustered at a household level.

Table 4: **Impact of Debt on ΔLog(Spending) Following Income Shocks**

	(1) OLS ΔLog(Spd)	(2) OLS ΔLog(Spd)	(3) OLS ΔLog(Spd)	(4) IV ΔLog(Spd)	(5) IV ΔLog(Spd)	(6) IV ΔLog(Spd)	(7) OLS ΔLog(Spd)	(8) OLS ΔLog(Spd)
ΔLog(Inc)	0.295***	0.258***	0.264***	0.377***	0.365***	0.336***	0.398***	0.474***
	(0.004)	(0.003)	(0.005)	(0.033)	(0.062)	(0.073)	(0.013)	(0.003)
ΔLog(Inc)*Debt/Assets		0.068***			0.077***			
		(0.003)			(0.027)			
ΔLog(Inc)*Debt/Income			0.060***			0.072**		
			(0.004)			(0.024)		
ΔLog(Inc)*Net Assets							-0.103***	-0.168***
							(0.006)	(0.002)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,600,000
Period FE	YES	YES	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES	YES	YES
Instrumented Variables	None	None	None	Inc	Inc,Lev	Inc,Lev	None	None
F-Tests (Inc)	-	-	-	45.9	36.1	42.6	-	-
F-Tests (Lev)	-	-	-	-	17.4	28.4	-	-
Source	Data	Data	Data	Data	Data	Data	Data	Simulation

*** p<0.01, ** p<0.05, * p<0.1

Notes: Columns (4)-(6) instrument for Δ Log(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer (Positive and Negative Earnings Surprises, Completed Acquisitions/Dispositions, Material Impairments, and Layoff Announcements) in the prior quarter as well as interactions of firm shocks and various leverage measures. In columns (5)-(6), the leverage term is instrumented with the Saiz (2010) measure of housing supply elasticity and housing supply elasticity interacted with homeowner status in 2008 in IV specifications. Debt to Assets measures total debt over total assets. Debt to Income measures the household's gross debt to annual income ratio. The dependent variable is ΔLog(Spending) by household-quarter. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a household level. Simulation run on 150,000 households over 24 quarters of simulated data. Parameters include a coefficient of risk aversion of 2, an annualized discount rate of 0.95. Income process mean and variance matched to data. Asset grid matched to data censored at 99th and 1st percentiles. Net Assets represent total household asset holdings (eg. property, equity, cash, bank account) minus total household debt (eg. mortgage, credit card, car loan, student loans) in millions of dollars.

Table 5: **Impact of Debt on ΔLog(Spending) Following Income Shocks - Robustness**

	(1)	(2)	(3)	(4)	(5)
	IV - No Interact	IV - MSA Trend	IV - State Banking	IV	IV
	ΔLog(Spd)	ΔLog(Spd)	ΔLog(Spd)	ΔLog(Spd)	ΔLog(Spd)
ΔLog(Inc)	0.389***	0.371***	0.418***	0.348***	0.381***
	(0.037)	(0.035)	(0.039)	(0.043)	(0.032)
ΔLog(Inc)*Debt/Assets	0.074***	0.069***	0.091*	0.055**	0.062***
	(0.031)	(0.025)	(0.046)	(0.026)	(0.024)
ΔLog(Inc)*Age				0.0081***	
				(0.001)	
ΔLog(Inc)*Age^2				-0.00012***	
				(0.000)	
ΔLog(Inc)*IncomeBin					-0.043***
					(0.011)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES
Instrumented Variables	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev
F-Tests (Inc)	33.9	42.2	31.8	37.8	38.5
F-Tests (Lev)	15.2	32.8	11.9	18.9	17.3

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: All columns instrument for ΔLog(Income) (and interactions with ΔLog(Income)) with Positive and Negative shocks to a household's employer (Positive and Negative Earnings Surprises, Completed Acquisitions/Dispositions, Material Impairments, and Layoff Announcements) in the prior quarter as well as interactions of firm shocks and various leverage measures. In column (1), the leverage term is constant across time, taken from the 2008 value, and is instrumented with the Saiz (2010) measure of housing supply elasticity. Column (2) features a time-varying leverage term and instruments for this term with the Saiz housing supply elasticity measure and this measure interacted with an MSA time trend. Debt to Assets measures total debt over total assets. Income Bin represents the self-reported level of household income in $25,000 bins. The dependent variable is ΔLog(Spending) by household-quarter. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a firm level.

Table 6: **Effects of Liquid Assets, Illiquid Assets, and Debt on ΔLog(Spending) Following Income Shocks**

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
Sample:	All	All	All	Non-Durables	Durables
ΔLog(Inc)	0.315***	0.343***	0.346***	0.319***	0.414***
	(0.031)	(0.026)	(0.023)	(0.022)	(0.021)
ΔLog(Inc)*(Debt/Inc)	0.076***	0.071***	0.051***	0.049***	0.063***
	(0.024)	(0.023)	(0.016)	(0.015)	(0.021)
ΔLog(Inc)*(Total Assets/Inc)		-0.049***			
		(0.014)			
ΔLog(Inc)*(Liq Assets/Inc)			-0.074***	-0.069***	-0.101***
			(0.014)	(0.016)	(0.018)
ΔLog(Inc)*(Non-Liq Assets/Inc)			-0.028***	-0.024**	-0.037***
			(0.010)	(0.011)	(0.015)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES
Instrumented Variables	Inc	Inc	Inc	Inc	Inc

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: All columns instrument for for ΔLog(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer in the prior quarter as well as interactions of firm shocks and leverage and credit measures. The dependent variable in columns 1-3 is Δ Log(Spending) by household-quarter. Columns 4 and 5 use a dependent variable of Δ log(Non-durables spending) and Δ log(Durables spending), respectively. Liquid assets measure savings accounts, checking accounts, cash holdings, and non-retirement equity accounts. Illiquid assets include pensions and retirement accounts, housing wealth, and other property holdings. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a firm level.

Table 7: **Impact of Debt and Credit Constraints on ΔLog(Spending) Following Income Shocks**

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
Sample:	All	All	All	All	All	All
ΔLog(Inc)	0.315***	0.343***	0.346***	0.329***	0.334***	0.324***
	(0.031)	(0.026)	(0.023)	(0.022)	(0.021)	(0.023)
ΔLog(Inc)*(Debt/Assets)	0.076***	0.069***	0.050***	0.033**	0.028*	0.022
	(0.024)	(0.023)	(0.017)	(0.015)	(0.016)	(0.015)
ΔLog(Inc)*(Credit Score)		-0.039***	-0.029**	-0.026**	-0.023*	-0.024**
		(0.014)	(0.011)	(0.012)	(0.013)	(0.012)
ΔLog(Inc)*(Unused Credit)			-0.058***	-0.057***	-0.059***	-0.053***
			(0.012)	(0.011)	(0.010)	(0.010)
ΔLog(Inc)*(Liq Assets)				-0.068***	-0.066***	-0.067***
				(0.015)	(0.016)	(0.019)
ΔLog(Inc)*(Credit Limit Decline)					0.057*	0.058*
					(0.033)	(0.031)
ΔLog(Inc)*(Marginal Int Rate)						0.090*
						(0.051)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Instrumented Variables	Inc	Inc	Inc	Inc	Inc	Inc

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: All columns instrument for for Δ Log(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer in the prior quarter as well as interactions of firm shocks and leverage and credit measures. The dependent variable is Δ Log(Spending) by household-quarter. All interaction terms are measured in 2008 and held constant over time for each household. Credit scores, Unused Credit, Liquid Assets, and Marginal Interest Rates are normalized to standard deviation of 1. Credit scores are self-reported FICA scores. Unused credit refers to room to borrow on existing credit cards and HELOCs. 'Credit Limit Decline' is an indicator that measures whether the household has seen a decline in a credit limit on one of their credit cards. 'Marginal interest rate' refers to the highest interest rate that a household has on an existing credit card, loan, or line of credit. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a firm level.

Table 8: **Impact of Debt on Spending Following Income Shocks - Positive and Negative Shocks**

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
Shock Type	All	Only Positive	Only Negative	All	Only Positive	Only Negative
	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)
Δ Log(Income)	0.336***	0.212***	0.441***	0.365***	0.177***	0.426***
	(0.029)	(0.015)	(0.012)	(0.022)	(0.018)	(0.019)
Δ Log(Income)*(Debt/Income)	0.072***	0.026	0.098***			
	(0.014)	(0.016)	(0.021)			
Δ Log(Income)*(Debt/Assets)				0.077***	0.024*	0.099***
				(0.027)	(0.018)	(0.029)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Instrumented Variables	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: Columns (1) - (6) instrument for Δ Log(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer in the prior quarter as well as interactions of firm shocks and various leverage measures. Columns (2) and (5) use only positive firm shocks. Columns (3) and (6) use only negative firm shocks. The leverage term is instrumented with the Saiz (2010) measure of housing supply elasticity. Debt to Assets measures total debt over total assets. Debt to Income measures the household's gross debt to annual income ratio. All standard errors clustered at a household level. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a firm level.

A Data Appendix

A.1 Housing Supply Elasticity

The second instrument utilized for this paper is one designed to identify exogenous variation in levels of debt and leverage, as well as identifying locations with larger exogenous changes in property prices. I use a measure of housing-supply elasticity developed by Albert Saiz (2011) that is based on geographic characteristics of metro areas. Saiz uses satellite-generated data on characteristics such as water boundaries, elevation, and elevation changes to construct an index of land unavailability covering the majority of the population of the United States. Saiz conducts his analysis at a CBSA level, providing an index for 257 of the 953 CBSAs, containing over 72% of the population of the United States as of 2010. All samples in this paper are restricted to be from households residing in zip codes that are contained within these CBSAs.

Glaeser, Gyourko, and Saiz (2008) as well as Mian and Sufi (2010), demonstrate the power of this instrument in both driving changes in housing prices and coincident changes in household leverage levels. They demonstrate the linkages between housing supply elasticity, home prices, and housing bubbles, both during the 1980s as well as during the early 2000s, in the period leading up to the Great Recession. In areas with relatively inelastic housing supplies, households homes' appreciated in value and home equity was often extracted, leading to higher levels of debt and leverage. Mian and Sufi (2013) demonstrate that housing supply elasticity did not systematically vary with things like local wage growth, exposure to or growth in housing-related sectors like construction, or large positive population growth. They find that more inelastic cities have higher per-capita income and per-capita net worth, but that these correlations are constant over time and thus, for the purposes of my paper, will not play a direct role in the differenced specifications.

Another concern with this variable is that housing supply elasticity acts as both a driver of household debt as well as a driver of asset prices. This is due to the nature of the housing boom and bust cycle during the late 2000s, where inelastic areas saw rapid increases in housing prices and housing-related debt followed by larger declines in asset prices. In my analysis of the impact of leverage on the responsiveness of consumption to income shocks, I use the variable to instrument solely for the initial level of leverage that the household faces. Concerns that the elasticity measure may also proxy for long-run decreases in housing prices may be mitigated by the fact that I use high-frequency changes in income while property prices decline only over a period of years. Moreover, the elasticity of local housing supply is uncorrelated with the firm shocks that drive changes household income, as most of these firms are large public, and geographically-spread, employers relatively unaffected by the decline in property prices. I also test specifications which include location level trends, property prices, or gross household assets as further controls for the decline in housing prices, finding no significant change in the magnitudes of my estimates.

A.2 Layoff Data

Data on the number of layoffs by firm-quarter are constructed from a search of newspaper articles regarding large instances of layoffs. The database in question is the Access World News Newsbank (AWN) database of national and local US newspapers. This database extends back to the 1980's and contains well over 1,500 newspapers during the same period of 2008-2013.

I first compile a complete list of full firm names (eg. 'Agilent Technologies Incorporated') and construct a complementary list of shortened firm names (eg. 'Agilent'). This shortening is largely done by removing all common terms and punctuation from firm titles. Using this list of full and shortened firm names, I search the AWN database for articles in which the title contains three terms. The first is the firm name (or shortened firm name). The title must also contain at least one term from each of the following sets: (positions OR jobs OR employees OR workers OR notices OR roles OR staff OR personnel) and (slash OR slashes OR slashing OR lost OR losses OR layoffs OR sheds OR axes OR cuts OR fires OR layoff OR shed OR axe OR cut OR cutting OR axing OR shedding OR reduce OR "lay off" OR "laid off"). Thus, I am attempting to identify news articles which are primarily about layoffs at a particular firm.

For each resultant article, in addition to the firm name in question, I save two additional pieces of information. The first is the date that the article was written. The second is the number of layoffs mentioned in the title. Full article titles regarding layoffs are often written in a formulaic manner (eg. "Wells Fargo to cut 4,000 jobs" or "Alcoa lays off 2,400 workers") and contain the number of layoffs that were announced by the firm. These searches yield approximately 16,000 news articles covering 512 firms from 1985 until 2013. I categorize and count the number of articles by firm-month, discarding any articles that do not have at least 2 articles regarding layoffs in that firm-month. Manual inspection reveals that these articles are generally false positives.

Finally, I use the most common number of layoffs, by firm-quarter, reported by newspaper articles as the true number of layoffs for that given firm-quarter instance of layoffs. Approximately 75% of firm-quarters have 75% of relevant articles agreeing on the number of layoffs in a given quarter. Just over 50% of firm-quarter observations have 100% agreement, with all articles regarding that layoff reporting the same number of layoffs. The final sample covers 179 firms and 330 instances of layoffs from

A.3 Job Switching

One potential worry in using negative firm shocks is that there are unlikely to be permanent or even semi-permanent in duration because individuals can switch jobs and negate the negative effects. I find evidence suggesting this is not the case; the income of users who switch jobs after a negative shock is insignificantly different than the income of users who switch jobs with no negative shocks.

Overall, users are not more likely to switch jobs after a positive shock, but are more likely to switch jobs after a negative shock. I test this with a probit of 3 lags of negative shock indicators, 3 leads of negative shock indicators, and contemporaneous negative shock indicators, finding significant effects on contemporaneous shocks of switching jobs in the following month. This remains true, though the coefficient decreases in magnitude, if only using non-layoff negative shocks (large and surprising negative earnings reports and write-offs by the firm).

I find that, on average income is depressed after leaving a job for the first three months. However, when transitioning to a new observable job (eg. to another employer from the NYSE or NASDAQ), income tends to increase in the months after leaving a job. Moreover, leaving a job after a shock does not significantly alter these relationships, with insignificant coefficients on interactions between negative shocks to firms and leaving that firm. This suggests that users cannot avoid the entirety of negative firm shocks simply by leaving their firm.

One explanation for this finding relies on users being rational members of the labor market and the presence of switching costs and firm-specific human capital exists. In this case, employees will not remain at a firm if there is an opportunity for a significant raise elsewhere. Moreover, on average, opportunities and pay at other firms will be lower than at the current firm due to the fact that firm-specific human capital has been accumulated. Thus, employees at firms hit by a negative shock will likely have no options to be employed at other firms at a significantly higher wage or else they would have already left.

Despite a lack of opportunities at higher wages, if household income was subject to large declines following a negative firm shock, employees may still be induced to change jobs to ameliorate sizable drops in income (eg. take another job at 10% less rather than endure a 30% decline in income). However, firm shocks in this paper cause drops in income on the order of 1-3%. Thus, considering a likely loss in income due to firm-specific human capital and job-switching costs, most users would be unlikely to be substantially better off if they switched jobs following a negative shock.

B Results Appendix

Figure B1: **Median Nominal Income Growth Expectations**

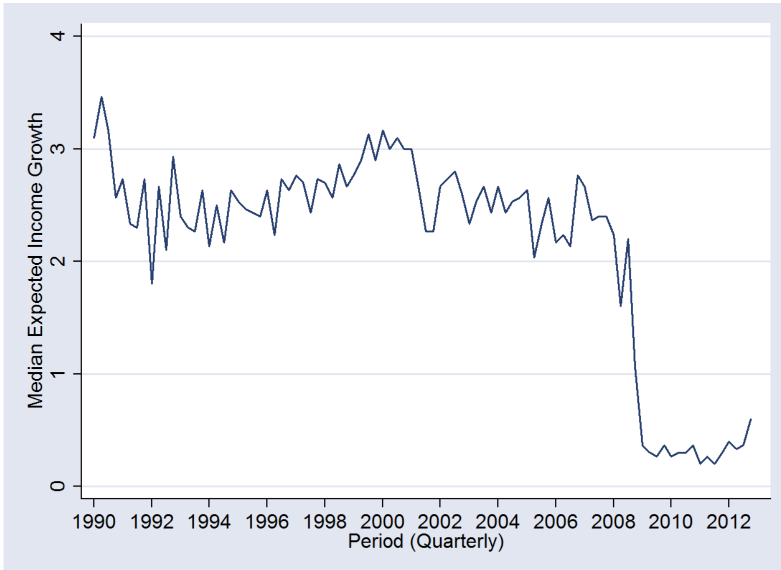

Notes: Median Nominal Income Growth Expectations, taken from data from the Michigan Consumer Survey. Data from the first quarter of 1990 to the fourth quarter of 2012. Data from survey answers to the question: “During the next 12 months, do you expect your (family) income to be higher or lower than during the past year?’ and ’By about what percent do you expect your (family) income to increase during the next 12 months?”

Table B1: **Summary Statistics Weighting Comparison**

	Mean	Median	CPS-Weighted Mean	CPS-Weighted Median
Male	0.72	1	0.65	1
Married	0.56	1	0.48	1
Homeowner	0.54	1	0.48	0
Income (Monthly)	$7,362	$5,678	$6,758	$4,810
Spending (Monthly)	$7,300	$5,646	$6,725	$4,790
Equity Wealth	$46,628	$3,391	$42,790	$902
Prop. Wealth (Homeowners)	$277,454	$220,666	$316,801	$232,500

Notes: CPS weighting is done at a household level based on age, sex, income, and state of residence. Households in personal financial website data are equated to head-of-households in the CPS data.

Table B2: **Layoffs by Firm**

Firm	Mon	Day	Year	Layoffs	Firm	Mon	Day	Year	Layoffs
AT&T	1	1	2010	160	MEDTRONIC	2	19	2011	2000
LOCKHEED	1	1	2010	1200	NORTHROP GRU.	2	26	2011	500
FOOT LOCKER	1	8	2010	120	WELLS FARGO	3	19	2011	400
NYSE	1	8	2010	350	CISCO	4	9	2011	550
UPS	1	8	2010	1800	LOCKHEED	6	11	2011	1200
HOME DEPOT	1	22	2010	1000	CISCO	7	9	2011	6500
VERIZON	1	22	2010	13000	BOSTON SCI.	7	23	2011	1400
WALMART	1	22	2010	11200	BANK OF AMER.	8	13	2011	3500
HUMANA	2	12	2010	1400	NORTHROP GRU.	8	20	2011	500
BOEING	2	19	2010	1000	BANK OF AMER.	9	10	2011	30000
IBM	2	26	2010	1518	LOCKHEED	9	24	2011	670
CHEVRON	3	5	2010	2000	LOWE'S	10	15	2011	1950
NETFLIX	5	14	2010	160	WHIRLPOOL	10	22	2011	5000
PFIZER	5	14	2010	6000	AMD	10	29	2011	1400
HEWLETT PACK.	5	28	2010	9000	ADOBE	11	5	2011	475
NATIONWIDE	7	2	2010	2070	MORGAN STAN.	12	10	2011	1600
WELLS FARGO	7	2	2010	3800	ARCHER D.M.	1	8	2012	1000
WINN DIXIE	7	23	2010	120	METLIFE	1	8	2012	2575
FEDEX	9	10	2010	1700	KRAFT	1	15	2012	1600
ABBOTT	9	17	2010	3000	ABBOTT	1	22	2012	700
AON	10	8	2010	1800	AMERICAN AIR.	1	29	2012	13000
NORTHROP GRU.	11	12	2010	380	MICROSOFT	1	29	2012	200
STATE STREET	11	26	2010	1400	NATIONWIDE	2	5	2012	625
AMERICAN EXP.	1	15	2011	550	SUPERVALU	2	5	2012	800
BOEING	1	15	2011	900	IBM	2	26	2012	1100
ABBOTT	1	22	2011	1900	GOODRICH	3	4	2012	500
PFIZER	1	29	2011	1100	LEVIS	3	4	2012	500

Notes: Layoffs derived from news-based analysis of layoff announcements. Articles from over 1,500 local, regional, and national newspapers from the Newsbank news archive service were utilized to determine the announcement date and the number of individuals laid off for each firm.

Table B3: **No Differential Effects of Firm Shocks**

	(1) OLS ΔLog(Inc)	(2) OLS ΔLog(Inc)	(3) OLS ΔLog(Inc)
Positive Firm Shock	0.01114***	0.00992***	0.00935***
	(0.00195)	(0.00181)	(0.00199)
Positive Firm Shock*(Debt/Assets)		0.00106	
		(0.0043)	
Positive Firm Shock*Housing Elasticity			-0.00091
			(0.0039)
Negative Firm Shock	-0.01389***	-0.01169***	-0.01241***
	(0.00240)	(0.00239)	(0.00278)
Negative Firm Shock*(Debt/Assets)		-0.00156	
		(0.00198)	
Negative Firm Shock*Housing Elasticity			0.00074
			(0.00142)
Observations	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES
Household FE	YES	YES	YES

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: Positive and Negative Earnings Surprises are indicators whether an earnings report differed from the mean expectation of earnings (from IBES data) by more than 2% of the stock price. Acq/Dis of Assets and Material Impairments are indicators derived from SEC 8K data. Layoffs are the number of layoffs derived from scraping AccessWorldNews newspaper database. NonFirm income is all non-paycheck and non-bonus income. Column (6) condenses the various firm shocks into 'Positive' and 'Negative' binary indicators, with Positive Earnings Surprises and Completed Acq/Disp classified as positive shocks and the rest as negative shocks. Regressions span January 2008-December 2013. All standard errors clustered at a household level.

Table B4: **Sensitivity to Own-Firm Performance**

	(1) Log(Spending)	(2) Log(Spending)	(3) Log(Spending)	(4) Log(Spending)
Monthly Firm Returns	0.00125	0.00169**	0.00127	0.00140
	(0.000934)	(0.000821)	(0.000935)	(0.000921)
Current*Firm Returns			0.00366***	0.00199***
			(0.000614)	(0.000545)
Observations	4,211,478	4,211,478	4,211,478	4,211,478
R^2	0.699	0.715	0.699	0.715
Period FE	YES	YES	YES	YES
Household FE	YES	YES	YES	YES
Wealth/Income Controls	NO	YES	NO	YES

*** p<0.01, ** p<0.05, * p<0.1

Notes: Monthly firm returns are cumulative monthly returns for the firm which an individual spent the most time working for in the sample for the entire period of employment in addition to 6 months prior to and 6 months subsequent to employment with that firm. Current*firm returns interacts these firm returns with an indicator denoting whether the individual was employed at the firm in that month. Wealth and income controls include logged values of total income and total wealth. Regressions span January 2008-December 2013. All standard errors clustered at a household level.

Table B5: **Impact of Debt on Spending Following Income Shocks - Components of Household Spending**

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	IV	IV	IV
	Δ Log(Spend)	Δ Log(Dur)	Δ Log(Debt)	Δ Log(Spend)	Δ Log(Dur)	Δ Log(Debt)
Δ Log(Income)	0.316***	0.422***	0.227***	0.353***	0.549***	0.241***
	(0.029)	(0.032)	(0.021)	(0.022)	(0.025)	(0.017)
Δ Log(Income)*(Debt/Assets)	0.068***	0.099***	0.023**	0.077***	0.103***	0.031***
	(0.014)	(0.017)	(0.011)	(0.027)	(0.033)	(0.012)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Instrumented Variables	-	-	-	Inc,Lev	Inc,Lev	Inc,Lev

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: Debt Spending represents spending on estimated credit card interest payments, loan payments, and mortgage payments. Durables spending is composed of auto payments, furnishings, auto service and parts, and home improvement spending. Columns (4) - (6) instrument for Δ Log(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer in the prior month as well as interactions of firm shocks and various leverage measures. The leverage term is instrumented with the Saiz (2010) measure of housing supply elasticity in IV specifications. Debt to Assets measures total debt over total assets. 'Positive' shocks are Positive Earnings Surprises, and Completed Acq/Disp while the remaining firm shocks (Negative Earnings Surprises, Material Impairments, and Layoff Announcements) are classified as 'Negative' shocks. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a household level.

Table B6: **Impact of Debt on Spending Following Income Shocks - Paychecks and Bonuses**

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
Shock Type	All	Only Paycheck	Only Bonus	All	Only Paycheck	Only Bonus
	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)	Δ Log(Spend)
Δ Log(Income)	0.336***	0.405***	0.161***	0.365***	0.427***	0.183***
	(0.029)	(0.)	(0.)	(0.022)	(0.)	(0.)
Δ Log(Income)*(Debt/Income)	0.060***	0.084***	0.021			
	(0.014)	(0.016)	(0.015)			
Δ Log(Income)*(Debt/Assets)				0.068***	0.087***	0.025
				(0.014)	(0.016)	(0.014)
Observations	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721	3,014,721
Period FE	YES	YES	YES	YES	YES	YES
Household FE	YES	YES	YES	YES	YES	YES
Instrumented Variables	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev	Inc,Lev

*** p<0.01, ** p<0.05, * p<0.1

Notes: Columns (2) and (5) use only changes in paycheck income as a dependent variable while columns (3) and (6) use only changes in bonus income as a dependent variable.. All columns instrument for Δ Log(Income) (and interactions with Δ Log(Income)) with Positive and Negative shocks to a household's employer in the prior month as well as interactions of firm shocks and various leverage measures. The leverage term is instrumented with the Saiz (2010) measure of housing supply elasticity in IV specifications. Debt to Assets measures total debt over total assets. 'Positive' shocks are Positive Earnings Surprises, and Completed Acq/Disp while the remaining firm shocks (Negative Earnings Surprises, Material Impairments, and Layoff Announcements) are classified as 'Negative' shocks. All regressions weighted by CPS-derived household frequency weights. Regressions span January 2008-December 2013. All standard errors clustered at a household level.