

Cooper, Daniel; Peek, Joe

Working Paper

The effects of changes in local-bank health on household consumption

Working Papers, No. 18-5

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Cooper, Daniel; Peek, Joe (2018) : The effects of changes in local-bank health on household consumption, Working Papers, No. 18-5, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/202921>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



The Effects of Changes in Local-Bank Health on Household Consumption

Daniel Cooper and Joe Peek

Abstract:

Focusing on localized measures of bank health and economic activity, and renters as well as homeowners, this paper uses an innovative approach to identifying households likely in need of credit to investigate the effect on household spending of a deterioration in local-bank health. The analysis shows that local-bank health tends to impact the expenditures of renters more than homeowners, with the strongest effects for households that likely need credit—those experiencing a negative income shock and having limited liquid wealth. These findings contribute to the discussion of the linkages between the financial sector and real economic activity.

JEL Classifications: E21, G21

Keywords: consumption, local-bank health, credit constraints

Daniel Cooper is a senior economist and policy advisor in the research department at the Federal Reserve Bank of Boston. His email address is daniel.cooper@bos.frb.org. Joe Peek is a vice president and economist in the Boston Fed's research department. His email address is joe.peek@bos.frb.org.

The authors thank María José Luengo-Prado, Jonathan Parker, and Bent Sørensen for extremely helpful comments, and they thank Mariya Sakharova and Katrina Truebenbach for excellent research assistance. They also thank Michael Corbett and Peggy Gilligan for their outstanding assistance constructing the bank health measures.

This paper presents preliminary analysis and results intended to stimulate discussion and critical comment. The views expressed herein are those of the authors and do not indicate concurrence by the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System.

This paper, which may be revised, is available on the website of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/wp/index.htm>.

This version: October, 2018

1 Introduction

The 2007–2009 US financial crisis and the associated Great Recession have stimulated volumes of research and discussion concerning linkages between the financial sector and real economic activity. With consumption making up roughly 70 percent of US Gross Domestic Product, the links between household expenditures and the financial markets are particularly relevant. The literature focuses primarily on wealth effects when studying the connections between consumption and the financial sector (for a review, see Cooper and Dynan 2016). Indeed, oft-cited papers analyzing the impact of the financial crisis on consumption concentrate on the role of plunging house prices or the related topic of household financial distress and deleveraging (see, for example, Mian, Rao, and Sufi 2013; Dynan 2012; Albuquerque and Krustev 2018). Although falling house prices and mortgage delinquency played important roles in the financial crisis, the linkages between the financial sector and household consumption are broader and more complex. Indeed, both the availability of credit and access to that credit are important determinants of household consumption. To finance and smooth their expenditures, many households, much like small businesses, rely on bank credit, the availability of which deteriorated during the financial crisis and more generally fluctuates over time in accordance with the health of the financial sector.

This paper explores the effect of local credit availability on household consumption beyond the standard income and wealth channels. In particular, we investigate the link between the health of banks operating in a locality (local-bank health) and the consumption of households in that locality using household-level data from the Panel Study of Income Dynamics (PSID). While banks help finance homeowners’ consumption through mortgages and home equity lines of credit, many households (including renters and especially those households that are liquidity constrained) rely on banks for personal loans, automobile loans, vacation loans, and other spending needs.¹ Moreover, to smooth consumption over time, households often use credit to purchase many durable goods in addition to autos, such as furniture and large appliances.² Thus, a deterioration in bank health that reduces banks’ ability and/or willingness to make credit available potentially affects a wide range of consumers. That is, the link between bank financing and household consumption

¹ Throughout the paper, we use the terms “liquidity constrained,” “borrowing constrained,” and “constrained” interchangeably to reference households that are more likely to need or want access to credit—especially when credit is less available—in order to remain on their desired consumption path.

² Of course, one can argue that the role of banks in providing direct consumer loans has diminished over time. Today, auto loans and loans for other durable goods frequently are provided by the captive finance arms of auto manufacturers or the retailers selling big-ticket household items, and the emergence of FinTech has resulted in rapid growth in the origination of personal loans by nonbanks. Still, the funding for such loans from captive finance firms and even some FinTech firms often comes in part from banks. Thus, bank financing likely continues to play an important role in consumer-credit availability, even though that role may be indirect.

goes well beyond an effect on homeowners operating through mortgage availability. While a decline in house prices contributes to a deterioration of bank health, and thus banks' ability and willingness to lend, the pullback in bank lending reduces the availability of consumer loans more generally, as well as the availability of mortgage and home equity loans.

Relative to homeowners, renters are more likely to need to borrow to smooth their spending. They tend to be younger than owners and have incomes that are both low relative to their expected future incomes and more variable than those of older, more established households. In addition, the impact on consumption depends not only on the availability of credit, but also on the ability of households to weather a period of reduced credit availability, for example, by relying on their liquid wealth.³ But in this respect, renters again tend to be at a disadvantage relative to homeowners, because renters often have less liquid wealth available to offset income shocks or expenditure shocks, such as large, unexpected medical bills.⁴ Thus, even though much of the literature focuses on homeowners and (declining) real estate wealth, it is possible that renters could be particularly sensitive to a reduction in credit availability emanating from a deterioration in bank health, even if the deterioration in bank health originated from a decline in house prices.

Households, whether homeowner or renter, that need or want to borrow to smooth their consumption path are likely particularly impacted by a contraction in credit availability. We use an innovative approach to help identify these (liquidity-constrained) households. In particular, we develop a method based on determining shortfalls in household current income relative to the level of income predicted by individual-specific age-earnings profiles. We further refine this indicator of potentially constrained households using data on individuals' ages and their liquid wealth holdings.

We find that a deterioration in local-bank health adversely affects household (food) consumption, even after we control for household income, wealth, and demographic characteristics, as well as (exogenous) measures of local economic activity. In particular, renters' consumption tends to be more responsive to changes in bank health than does owners' consumption. More direct analysis of likely liquidity-constrained households confirms that fluctuations in bank health generally matter more for households that potentially need to borrow and whose spending therefore is most sensitive to credit availability.

We obtain these results from estimating a standard consumption function augmented with a measure

³ In fact, Damar, Gropp, and Mordel (2014), using Canadian household-level data to examine the relationship between bank health and consumption, find that a reduction in household credit supply reduces consumption for liquidity-constrained households, but not for households with sufficient liquid wealth that can finance their spending by drawing down their savings as needed.

⁴ And it is not just renters who are struggling with limited liquid wealth. For example, even after a period of extended economic recovery, the Federal Reserve's "Report on the Economic Well-Being of U.S. Households in 2017" finds that, "Four in 10 adults, if faced with an unexpected expense of \$400, would either not be able to cover it or would cover it by selling something or borrowing money. This is an improvement from half of adults in 2013 being ill-prepared for such an expense." (Board of Governors of the Federal Reserve System 2018)

of local-bank health. We exploit both cross-sectional and time-series variation in bank health at the local level—defined as a metropolitan statistical area (MSA) or the rural (non-MSA) portion of a state—to identify the effect of bank-loan supply on consumption. Our sample includes data from 1985 through 2015—a time period that includes both the banking crisis of the late 1980s and early 1990s and the more recent financial crisis associated with the Great Recession. Both episodes were characterized by major problems in real estate markets that severely damaged bank health; commercial real estate played a key role in the banking crisis, and residential real estate problems were central to the financial crisis. Consequently, our primary local-bank health measure is based on nonperforming real estate loans issued by banks with local branches. We are careful to identify local banking options available to a household by focusing on a bank’s branch locations rather than the location of a bank’s headquarters. Whereas the largest banks dominate the first-mortgage market, local banks and local branches of larger banks likely play a more important role in providing home equity and consumer loans.⁵

Our findings are robust to different approaches to measuring the health of the US financial sector and to employing broader measures of household consumption. Alternative bank health gauges include using bank supervisory data as well as a measure based on multi-locational banks with local branches that is arguably exogenous to local economic conditions. Multi-locational banks are those with one or more branches in a given location that account for a small enough share of the bank’s total deposits for the bank’s health to be considered exogenous relative to the location’s economic health. The broader consumption measures include one based on additional expenditure data available in the PSID starting in 1999 as well as one derived from the imputation approach in Blundell, Pistaferri, and Preston (2006).

This paper encompasses the literature on how income, wealth, and liquidity constraints affect consumption (see, for example, Lehnert 2004; Hurst and Stafford 2004; Johnson, Parker, and Souleles 2006; Cooper 2013; Campbell and Cocco 2007; Jensen and Johannesen 2017), along with the literature showing that bank-loan supply shocks (appropriately measured) can impact the real economy (see, for example, Chava and Purnanandam 2011; Bassett et al. 2014). Our results provide a quantitative framework for thinking about the impact on consumption of credit supply shocks from a banking crisis or more general financial crisis—shocks not necessarily aligned with the aggregate business cycle.

Other related papers include Abdallah and Lastrapes (2012), which examines the response of retail

⁵ In the context of small-business lending, Nguyen (Forthcoming), among others, shows that even after the improvements in information technology, distance from bank branches still plays an important role in access to credit in the 2000s—especially during periods of tighter lending standards. One would expect the same to be true for personal loans to households. In fact, in the context of auto loans, Argyle, Nadauld, and Palmer (2017) find that the median borrower who obtains an auto loan from a bank uses a branch that is within a 15-minute drive from his or her home.

spending to a reduction in credit supply constraints in Texas; Greenstone, Mas, and Nguyen (2014), which studies how credit shocks impact the real economy; Agarwal and Qian (2017), which looks at the effects on consumption of a borrowing-related policy change in Singapore that reduced homeowners’ ability to obtain credit against their housing wealth; and Agarwal et al. (2018), which examines households’ propensity to borrow when their credit limits increase as well as banks’ marginal propensity to lend when their funding costs fall. Relative to these papers, our research is generally broader; we are interested in the impact of changes in the availability of local-bank credit on household spending by both homeowners and renters, highlighting the effects on liquidity-constrained households.

The remainder of the paper proceeds as follows. Section 2 discusses our data and how we construct our measures of bank health. Section 3 outlines our empirical approach. Section 4 presents our approach for identifying constrained households. Section 5 presents our main results, and Section 6 considers the robustness of our findings. Section 7 concludes.

2 Data

2.1 Household Data

Our analysis relies on household-level data from the 1984 through 2015 waves of the PSID, which is a longitudinal survey that began in 1968, was conducted annually through 1997, and has been conducted biennially since 1997. The PSID follows the original (1968) households and their offspring over time.

In each wave, the PSID contains detailed demographic, household income, and homeownership data and, if applicable, data for households’ self-reported home values and any outstanding mortgage debt (first and second liens). It also includes data on households’ food consumption—both at home and away from home—which serves as our baseline measure of household expenditures.⁶ Financial wealth and non-housing debt data are available in the PSID surveys for 1984, 1989, and 1994, and from 1999 onward. While most variables capture household information as of the survey year, income and some other labor market variables cover the year prior to the survey year. Section A.1.1 in the Appendix provides additional relevant details about the PSID data.

We chose the PSID for this study and not the Consumer Expenditure Survey (CEX) or another household-level data set for a few reasons. First, the longitudinal data allow us to track the spending behavior of a given

⁶ The PSID tracks households’ out-of-pocket expenditures on food, as well as any expenditures made using food stamps. We combine these two sources of food expenditures in our analysis. Food expenditure data were not collected in the 1988 or 1989 PSID waves.

household over time and control for household-specific spending habits. The PSID also contains comprehensive data on household income and wealth as well as data on household (food) expenditures that, while limited in scope, are available consistently for an extended period of time. Also, having access to detailed information in the PSID about where households live (through a restricted data contract)—information that is not as readily available with a dataset such as the CEX—is important for our analysis. These location data enable us to connect a household to the health of its local bank branches and to relevant local economic conditions.

2.2 Bank Data

Our baseline measure of bank health combines balance sheet information from banks' Consolidated Reports of Condition and Income (Call Reports) with individual bank-branch deposit data from the FDIC's Summary of Deposits (SOD). We use these data to generate measures of bank health at the local (MSA) and state non-MSA levels. Given the major consolidation of the banking sector over time, a bank's headquarters location is typically no longer representative of where it does most of its business, except for the smallest banks. This makes the SOD data an important component for calculating our measures of the health of banks that provide local services to households. The variables used for constructing bank health measures are reported at the consolidated bank (parent-bank) level rather than at the bank branch level. Thus, we construct our bank health measures by assigning each branch the value of the bank health measure of its parent bank for that period.⁷

Defining Local Areas (Geographic Locations) for our Analysis

Because the PSID does not contain information on the identity of the specific bank(s) from which a household obtains credit, we instead link households with the set of banks that operate in their local market.⁸ Under a restricted data contract designed to protect the confidentiality of the respondents, we obtained detailed geographic information about the households in the PSID. These data pinpoint the location of PSID households in each wave down to the census tract level. Because metropolitan areas are likely representative of the local banking markets for urban households, we use the available core-based statistical area (CBSA, commonly known as MSA) geocode data to assign each PSID household to a specific geographic location. If a household in a given state resides outside an MSA (or lives in an MSA that has too few banks as

⁷ In addition, it makes sense to examine bank-level health and not branch-level health because internal capital markets operate among branches within a bank (see Berrospide, Black, and Keeton 2016; Cortés and Strahan 2015).

⁸ Our underlying assumption is that households rely on local bank branches for most consumer credit and home equity loan needs, unlike the more national market for first mortgages and credit cards.

described below), we assign it to the state’s rural (non-MSA) area when linking households with local banking conditions.

We use the annual SOD deposit data, measured as of June 30 of each year, to identify all commercial and savings bank branches located in every (PSID) MSA or state rural area in each year. Because we combine annual SOD information with quarterly Call Report data to determine local deposit shares, we convert the SOD data to a quarterly frequency by using the annual SOD deposit value for each quarter in that calendar year. To account for changes in bank composition from quarter to quarter in a given location, for example due to closures or acquisitions—changes not reflected in the annual SOD data—we restrict our analysis to branches of parent banks that have Call Report data for a given quarter.

Our primary measure of bank health is based on the distribution of deposits and nonperforming real estate loans across banks with branches in a given location. As a result, we need enough individual banks to adequately isolate the lower tail of the relevant distribution by location. Consequently, we exclude MSAs that have fewer than five banks at any time during our sample period.⁹ Since not all PSID households live in an MSA or are located in an MSA that meets our minimum number of banks criterion, we aggregate the non-MSA locations to construct one “rural” (non-MSA) bank health measure by state and year.¹⁰

Bank Balance Sheet (Call Report) Data

We use quarterly Call Report data to construct indicators of bank health based on the share of nonperforming real estate loans on banks’ balance sheets. A key part of the data cleaning and variable construction process involves selecting a relevant sample of banks. In particular, we focus our analysis on Federal Deposit Insurance Corporation (FDIC)–insured commercial and savings banks headquartered in the 50 US states and the District of Columbia. We discuss additional restrictions we make in generating our sample of relevant banks in the Appendix (see Section A.1.2).

Because we want to capture local-bank health for areas with multiple banks, we compute a location-specific measure based on the health of individual banks operating in that location. Rather than equally weighting these bank branches, we use deposit-weighted measures, arguing that the relative importance of a bank to a given location is tied to its share of the total deposits of that location’s branches.¹¹

Credit availability is most important for households that rely on bank financing to smooth their consump-

⁹ To deal with MSAs where the number of unique banks fluctuates over time between fewer than five and five or more, our minimum-bank rule retains MSAs where there are at least five banks 50 percent of the time, but only for those subperiods when the location has a span of four-plus years with five or more banks.

¹⁰ New Jersey, Delaware, Rhode Island, and Washington, DC, do not have rural areas, given our definitions.

¹¹ Weighting based on a branch’s loan originations in the local market is preferable but not feasible, given the data availability.

tion expenditures over time, particularly liquidity-constrained households. Thus, we focus on an indicator of a curtailment of credit availability: the lower tail of the bank health distribution. A bank under severe financial stress—for example, due to current and prospective losses associated with a sharp increase in its nonperforming loans—may curtail its lending as it tries to improve the health of its balance sheet. A healthy bank that is not facing any binding (or near-binding) regulatory constraints, on the other hand, is likely to adjust its lending only gradually in response to changes in its health.

Our baseline measure of bank health is based on a nonperforming real estate loan ratio (RENPL), defined as real estate loans past due 90 days or longer and still accruing, plus nonaccrual real estate loans, divided by total bank assets. We focus on the health of the bank at the 10th percentile of the RENPL distribution in each location. We identify this bank by ordering the bank branches in each period and location by their parent bank’s RENPL values, from poorest health (highest RENPL value) to best health (lowest RENPL value). We then cumulate local bank deposits, starting with the branch (bank) that has the highest RENPL ratio (worst health). The RENPL value for the bank at the 10th percentile of the local deposit distribution (RENPL10) is our measure of bank health for that location and period.¹²

We base our bank health measure on nonperforming real estate loans rather than on total nonperforming loans because the two major episodes of banking stress during our sample period were real estate related: The banking crisis of the late 1980s and early 1990s was tied to commercial real estate problems, and the 2007–2009 financial crisis was tied to a residential real estate price collapse.¹³ Moreover, we prefer basing our bank health measure on a nonperforming loan ratio instead of another bank balance sheet measure, such as a capital-to-assets ratio, because, for example, capital-to-asset ratios exhibit systematic differences across bank asset-size classes, and a bank’s willingness to lend is likely more closely tied to the difference between its desired and actual capital ratios rather than its actual capital ratio. In addition, nonperforming loan ratios exhibit greater variation, improving the chance of identifying the bank health effects in our analysis.

Although we control for local economic activity in our analysis (see Section 2.3), the potential endogeneity of our bank health measure remains a concern, given that the health of the local economy likely affects the health of banks operating locally. This is certainly true for small banks for which a particular location accounts for a majority of their operations. However, the strength of economic activity occurring

¹² For example, assume that a given location has total bank-branch deposits of \$100 million. First, we order the bank branches in that location by the RENPL of their associated parent bank. We then start with the bank branch with the worst RENPL and cumulate deposits until we reach a total of \$10 million. The bank branch accounting for that cumulated deposit level is identified as the bank branch at the 10th percentile of the RENPL distribution, and we record the RENPL value of that branch’s parent bank for that location, year and quarter as the value of our RENPL10 measure.

¹³ An analogously constructed measure of total nonperforming loans, instead of nonperforming real estate loans, yields similar results, which is expected given that the two measures track each other closely.

in a particular location will have little measurable effect on the health of a multi-locational bank (ML BH) if that location accounts for very little of its operations. Because the health of such multi-locational banks can be considered exogenous to that location, we include multi-locational bank health measures as part of our robustness tests. Section A.1.3 in the Appendix provides details about how we calculate ML BH.

Bank Supervisory Data

As another robustness check, we also measure bank health based on supervisory data instead of banks' balance sheet information. Bank supervisors produce confidential bank ratings, known as CAMELS ratings, with integer ratings from 1 (best) to 5 (worst). A CAMELS rating is based on a bank's capital adequacy, asset quality, management, earnings, liquidity, and sensitivity to market risk. We construct a weighted average of the CAMELS ratings of banks with one or more branches in a given location, using bank branches' deposit shares (from the SOD) in the local market as the weights.¹⁴ Unlike with our balance sheet measure (RENPL10), we do not use a tail-based measure of the CAMELS ratings distribution because CAMELS ratings are discrete, and during most time periods relatively few, if any, banks in a given location have the lowest (4 or 5) ratings. A tail-based measure of bank health, therefore, would have a predominant value of 3, with little variation across time and location. Similarly, a measure based on the share of deposits in a location associated with a CAMELS 4- or 5-rated bank would be dominated by observations with a value of zero.

2.3 Local Economic Data

To help isolate the effects of bank health on consumption, our estimates control for local economic conditions—either at the MSA level or at the state level for households residing in non-metro areas—which might be correlated with the health of banks with local operations. Given our focus on nonperforming real estate loans, one control is real house price growth (HPG). Since HPG is location specific, and not tied to homeownership status, it can control for local economic activity for both renters and homeowners. The house price data come from CoreLogic, a private company that constructs constant-quality price indices at, among others, the CBSA (MSA) and state levels. We discuss how we construct HPG using these data in the Appendix (see Section A.1.4).

We also condition our estimates on local business cycle conditions, so that our bank health measure

¹⁴ Often, there is a delay between when a new bank is formed and when it receives its first exam and CAMELS rating from bank supervisors. In such a situation, we assign the CAMELS rating from the first exam to all dates the bank files a Call Report prior to its first supervisory exam, keeping in mind that the first two years of a newly formed bank have already been omitted.

does not serve as a proxy for the effect of fluctuations in the local economy. While employment growth is often used as a local business cycle control, bank health and employment growth are likely correlated, given that local loan supply presumably impacts local employment. Therefore, we employ a Bartik-style approach to account for local labor market conditions (Bartik 1991). In particular, we combine data on national employment growth by industry with data on industry composition at the local level to generate local “Bartik employment growth” (BEG). Since industry composition is slow to change, and thus likely uncorrelated with local bank health or consumption, this approach provides an arguably exogenous measure of employment growth by location. As with HPG, we calculate BEG at both the MSA and state level, assigning state-level BEG to rural (non-MSA) PSID households. See Appendix Section A.1.5 for further details on how we construct BEG.

2.4 External Validity: Bank Health

We check the external validity of our baseline measure of bank health by examining its predictive power for (real) local nonmortgage loan growth. In particular, we calculate real per capita nonmortgage loan growth at the MSA level using a 5 percent random sample of US credit reports from Equifax provided by the Federal Reserve Bank of New York Consumer Credit Panel (CCP). Including only those MSAs that have households in the PSID, we regress annual real loan growth on our (one-year) lagged baseline measure of bank health (RENPL10), as well as on controls for lagged local economic conditions (HPG and BEG). We also include year and location fixed effects. Appendix Section A.3 provides further details about the Equifax data and this analysis.

The regressions show a strong negative correlation between our measure of bank health and real per capita nonmortgage loan growth (see Table 1, column 1), all else being equal. Recall that high levels of RENPL10 indicate poor bank health, so the results are consistent with real loan growth declining as bank health deteriorates. The negative relationship is even stronger if we include home equity loans (but not primary mortgages) in our measure of loan growth (column 3). In terms of magnitudes, going from the 25th to 75th percentiles of the RENPL10 distribution (2.13 percentage points) leads to a decline in loan growth of 0.4 percentage point (column 1) to 0.9 percentage point (column 3), depending on the specification. In addition, to avoid concerns about local economic conditions affecting bank health and also loan growth, we regress real per capita loan growth on our ML BH measure (columns 2 and 4), which is arguably exogenous with respect to local economic conditions. Again, we find a strong, negative relationship between bank health and local loan growth. Therefore, it appears that our measure of bank health is relevant for capturing

local credit availability.

3 Estimation Approach

3.1 Empirical Specification

Our baseline approach projects log real household (food) consumption expenditures on log real household income, bank health, and a series of financial, demographic, and local economic controls. The estimates also include household-specific fixed effects to account for any time-invariant differences in spending behavior across households, as well as time fixed effects to capture broad macroeconomic trends. More specifically, we estimate regressions of the following form:

$$c_t^{i,j} = \beta_0 + \beta_1^k y_{t-1}^{i,j} + \beta_2^k BH_{\hat{t}-1}^j + \beta_3^k HPG_{\hat{t},\hat{t}-1}^j + \beta_4^k BEG_{\hat{t},\hat{t}-1}^j + \beta_5^k HY_{t-1}^{i,j} + \beta_6^k FY_{t-1}^{i,j} + \gamma^k \mathbf{X}_t^{i,j} + \delta_t^k + \eta^i + \epsilon_t^{i,j}, \quad (1)$$

where $c_t^{i,j}$ is the logarithm of annual real consumption for household i living in location j at time t , and $y_{t-1}^{i,j}$ is real after-tax family income for household i in location j for the previous year. $BH_{\hat{t}-1}^j$ is our appropriately timed measure of bank health in location j . $HPG_{\hat{t},\hat{t}-1}^j$ and $BEG_{\hat{t},\hat{t}-1}^j$ are appropriately timed house price growth and (Bartik) employment growth in location j . $HY_{t-1}^{i,j}$ is the lagged ratio of household housing wealth to income. $FY_{t-1}^{i,j}$ is the lagged ratio of household financial wealth to income.¹⁵ $\mathbf{X}_t^{i,j}$ is a vector of household demographic controls, including head age and age squared, family size, family size squared, homeownership status, and whether the household includes multiple earners.¹⁶ δ_t^k is a set of annual dummy variables (year fixed effects). η^i is a household-specific fixed effect.

This empirical approach is based on the hybrid log-level consumption function proposed in Muellbauer (2008) and used in Cooper (2013). This levels-based setup with the PSID data allows us to include households that have zero or negative wealth, and avoids having to combine two-year growth rates based on the biennial data starting in 1999 with annual growth rates through 1997.

Equation (1) pools renters and owners but allows the estimated coefficients for the two groups to differ, as captured by the superscript k on the parameters where $k = 0$ for renters and $k = 1$ for owners. Households may switch from being owners to renters during our estimation period, subject to the sample restrictions

¹⁵ Since financial wealth data are not available in every PSID wave, our sample size is reduced by about half. Consequently, we also report estimates of equation (1) that exclude $FY_{t-1}^{i,j}$ in order to make use of the larger sample.

¹⁶ We include a (1, 0) dummy variable for whether a household includes a spouse/cohabitant but at most one earner, and a (1, 0) dummy variable for whether a household includes a spouse/cohabitant and dual earners. The third category, single, is the excluded category.

outlined in Section 3.3. We calculate and report absolute effects for renters and owners for ease of comparison and interpretation. An alternative approach would be to split the sample by homeownership status. However, by pooling owners and renters, we maintain a larger sample size (due to the need for at least two observations to estimate a household’s fixed effect) and potentially increase the power of our estimates. In addition, this specification restricts each household’s fixed effect (η^i) from changing at any time during the sample period, even if the household switches between renting and owning its home.

Our coefficient of interest is β_2^k . Estimates of β_2^k that are different from zero imply that bank health has an independent effect on consumption for a given household type, even after controlling for local economic conditions and other demographic and financial factors that are typically important predictors of household spending. Recall that our primary (baseline) measure of bank health, RENPL10, is constructed such that higher values represent banks that are in worse health. As a result, finding $\beta_2^k < 0$ suggests that consumption falls when bank health deteriorates and rises when bank health improves. We further anticipate that the estimated bank health effects will be stronger (in absolute value) for constrained households that need credit to help smooth their consumption over time.

3.2 Data Timing

We align the data used to estimate equation (1) to account for local economic conditions, bank health, and other components of households’ information sets at approximately the time a household makes its consumption decision, given the timing of the available data in the PSID. Variables dated t in equation (1) are measured at the time of each household’s PSID interview, while variables dated $t - 1$ cover the year prior to the survey year. For example, income data, which cover the calendar year prior to the survey year, are dated $t - 1$ following this timing convention. Household wealth data, and hence the housing wealth-to-income or financial wealth-to-income ratios, are also measured as of $t - 1$ to capture a household’s stock of assets prior to its consumption decision. That is, to predict (food) consumption in 1985, we measure household wealth and income as of 1984. Due to data availability, household wealth data are lagged two calendar years when constructing wealth-to-income ratios once the PSID switches to biennial data in 1999.¹⁷ Including financial wealth limits our sample, because, given our timing conventions, these data are available prior to 1999 only in the 1985, 1990, and 1995 sample years. Demographic variables, such as age, family size, and homeownership, are measured at the time of the survey (t). Single versus dual-earner households

¹⁷ Given this timing of the survey data, housing wealth and financial wealth are scaled by average household income from the years on either side of the year wealth is measured. For example, (housing or financial) wealth in 2009, which we pair with 2011 consumption, is scaled by average household income from 2008 and 2010. Averaging the 2008 and 2010 values approximates income in 2009, which is not reported in the PSID.

are determined primarily based on data for the year prior to the survey year ($t - 1$). Since we identify constrained households using earnings information, the constrained indicator is also based on data from the year prior to the survey year.

In addition, we incorporate our bank health and local economic conditions data relative to the year and quarter in which a household is interviewed in each wave. We include (four-quarter) HPG and (four-quarter) BEG lagged one quarter relative to a household’s interview quarter (denoted by \hat{t} in equation [1]). Bank health is lagged one year (four quarters) relative to a household’s interview quarter (denoted by $\hat{t} - 1$). We employ this timing convention because, roughly speaking, it controls for bank health at the beginning of the period in which households make annualized spending decisions. Section A.1.6 includes a concrete example of how we deal with data timing in our estimation.

3.3 Estimation Sample

As with other PSID research, our sample and analysis are based primarily on the characteristics of the household head. Specific details of our sample construction (and restrictions) can be found in Section A.1.7 of the Appendix. Generally speaking, the sample runs from 1985 through 2015 and includes households in which either the head or spouse/cohabitant (if any) is at least 18 years old and no older than 64.¹⁸ Households in which a spouse/cohabitant who meets the age criteria is present but both the head and spouse/cohabitant report being students or retirees are also dropped, as are single (non-cohabitating) households in which the head reports being a student or a retiree. We also exclude households that report any business income, since the PSID imputes business owners’ labor earnings, which tend to be volatile and are often negative.

To be included in our sample, households must also have non-missing data for all demographic variables included in equation (1), as well as non-missing wealth data (housing wealth and financial wealth, depending on the specification). Further, we exclude households added to the PSID temporarily in the 1990s as part of the Latino or immigrant subsamples. Households with missing or negative family income data are also dropped, as are households that report zero family income or zero food expenditures. Finally, as discussed in more detail in the Appendix, we remove or winsorize outliers to prevent them from exerting undue influence on our results. Our final regression includes 39,888 renter observations and 53,188 owner observations. The sample includes 14,020 unique households.

¹⁸ Households are dropped only in the year(s) that they do not meet the sample criteria.

4 Identifying Constrained Households

Overview

Whereas the Survey of Consumer Finances (SCF) contains direct questions about households' need for and access to credit, most datasets, including the PSID, do not, and information on whether households are liquidity or credit constrained must be inferred. Indeed, some papers, such as Johnson and Li (2010), have used the relationship between answers to the credit questions and household observables in the SCF to infer credit needs in the PSID and other datasets. However, without direct questions, there is not necessarily a one-size-fits-all approach to determining whether a household is constrained, and imputation approaches, especially across datasets, likely misclassifies some households. Identifying constrained households based on the age of the head of the household (see, for example, Lehnert 2004) or liquid asset holdings (see, for example, Zeldes 1989) are also reasonable approaches but still likely to misclassify some households.¹⁹

We take, as an alternative, a comprehensive and innovative approach that uses information on households' income shortfalls in a given period to identify those households with a likely need to obtain credit to smooth their consumption. We identify these shortfalls based on households' current income relative to the level of income predicted by age-earnings profiles with individual-specific intercepts. This approach allows us to identify years in which households face negative income shocks.²⁰

Estimating Households' Predicted Income: Age Earnings Profiles

Households' predicted income is derived from first estimating standard age-earnings profiles in the PSID by gender and education group:

$$y_t^{\omega,s} = \beta_0 + \beta_1^{\omega,s} age_t + \beta_2^{\omega,s} age_t^2 + \epsilon_t^{\omega,s}, \quad (2)$$

where $y_t^{\omega,s}$ is labor earnings at time t for an individual of a given gender (s) and education level (ω).

We estimate equation (2) at the individual level using all the available earnings information for the head of the household and any spouse/cohabitant separately. Individuals are assigned to one of four education

¹⁹ For example, a household headed by a relatively young person could have relatively high income, either absolutely or relative to its future income, and live within its current means with little desire to obtain credit from banks or other lenders. In addition, so-called wealthy hand-to-mouth households accumulate debt for purchases other than their primary residence even when they have adequate liquid wealth as a hedge against being more constrained in the future (see Kaplan, Violante, and Weidner 2014).

²⁰ Generating and evaluating individuals' age-earnings profiles in datasets such as the PSID is not new; however, these profiles are almost always used to evaluate income variability and/or uncertainty (see, for example, Kazarosian 1997; Charles and Hurst 2003; Hurst et al. 2010). We extend the existing approaches to incorporate individual-specific intercepts, and take an innovative approach by using these age-earnings profiles to evaluate whether individuals are constrained.

groups—less than high school, high school, some college, and college or more—and we estimate age-earnings profiles separately for each of these groups by gender.²¹ Since household-level data are inherently noisy and individuals potentially have years when they do not work or they work a reduced amount of time (due to, for example, unemployment, retirement, being in school), we have a number of procedures for cleaning the data to generate more accurate estimates of $\beta_1^{\omega,s}$ and $\beta_2^{\omega,s}$ from equation (2) (see Section A.2 in the Appendix for further details).

We use our estimates $\hat{\beta}_1^{\omega,s}$ and $\hat{\beta}_2^{\omega,s}$ (by gender and education group) to determine individual-specific age-earnings profiles. These profiles include individual-specific intercepts, which account for the fact that even though the slope of the age-earnings profiles may be similar, for example, for two college-educated workers of the same gender, their intercepts may differ reflecting the careers (industries) the individuals choose. The individual-specific age-earnings profile intercept, α_i , is constructed for each observation as:

$$\alpha_i = y_{i,j}^{\omega,s} - \hat{\beta}_1^{\omega,s} age_j - \hat{\beta}_2^{\omega,s} age_j^2, \quad (3)$$

where $y_{i,j}^{\omega,s}$ is the labor earnings of individual (i) at age (j) who is in education group (ω) and gender group (s). Thus, each age and income combination for a given individual will potentially yield a different estimate of α_i . We use the median value, $\bar{\alpha}_i$, of these estimates to construct the individual's predicted earnings, $\hat{y}_{i,j}^{\omega,s}$, at each age j that he or she is in the labor force.

$$\hat{y}_{i,j}^{\omega,s} = \bar{\alpha}_i + \hat{\beta}_1^{\omega,s} age_j + \hat{\beta}_2^{\omega,s} age_j^2. \quad (4)$$

To improve accuracy, we calculate $\bar{\alpha}_i$, and hence an individual's predicted earnings, only if we have at least five earnings (and age) observations for the individual.²²

Our variable of interest for determining likely constrained individuals is $dev_{i,j}$ —the percentage deviation of an individual's current labor earnings from his or her (individual-specific) predicted labor earnings at a given age j :

$$dev_{i,j} = \frac{y_{i,j}^{\omega,s} - \hat{y}_{i,j}^{\omega,s}}{\hat{y}_{i,j}^{\omega,s}}. \quad (5)$$

This variable indicates a whether an individual has a negative deviation in his/her labor-earnings ($dev_{i,j} < 0$)

²¹ We assign individuals to the education group consistent with their highest observed education level.

²² Observations where the individual is unemployed, retired, not in the labor force, a business owner, or a student do not count toward the minimum of five. However, we include observations where the individual is unemployed if we have at least five other observations where the individual has labor income.

or a nonnegative deviation ($dev_{i,j} \geq 0$). For dual-earner households, we combine labor earnings for the head and spouse/cohabitant as well as their individual predicted earnings to calculate (negative or nonnegative) income deviations at the household level.

Determining Likely Constrained Households

We classify potentially constrained households two ways: first as any household with a negative income deviation in a given period, and second as households with negative income deviations of 5 percent or more. In both cases, unconstrained households are those with nonnegative income deviations. In some specifications, we further refine our definition of constrained households by combining information on negative income deviations with data on the age of the earner(s) in the household or with the households' relative amount of liquid wealth (above/below median). Even in the absence of a negative income deviation, households that are younger are more likely to demand credit in an effort to raise current consumption toward their lifetime average. And households with lower liquid wealth holdings that experience negative income shocks have less ability to use savings to maintain their desired consumption path and are thus more likely to demand credit in order to maintain their desired consumption path.

5 Empirical Results

Overview

Our analysis of how the effects of bank health on consumption differ between renters and owners potentially would be problematic if renters and owners are distributed across locations in a way that they face different economic conditions or bank health. The statistics in Table 2 show that this is not the case. Indeed, the local economic conditions (HPG and BEG) faced by renters and owners are quite similar. In addition, and importantly, there are no meaningful differences between the locations where homeowners live and where renters live in terms of our baseline (RENPL10) and supervisory (CAMELS) measures of local-bank health. As a result, any observed differential effect of bank health on consumption for renters as compared with owners is unlikely to be driven by one group experiencing better (or worse) local-bank financing conditions or other local economic conditions on average.

The statistics in Table 2 further show that renters have lower average real food consumption than homeowners, consistent with renters' lower reported average real after-tax income levels as well as the fact that they tend to be younger and have smaller families. In addition, homeowners tend to have higher financial-

wealth-to-income ratios than renters, and there is a greater share of dual-earner homeowner households than dual-earner renter households. Consistent with our prior that renters are more likely to be constrained, we find a slightly greater share of renters who are constrained, on average, than owners based on our income deviation definition. This is unsurprising insofar as one might expect negative income shocks to be spread evenly across owners and renters except to the extent that renters tend to be younger and have, on average, lower education and incomes—factors that tend to lead to somewhat more spells of unemployment. Importantly, should households suffer negative income shocks, renters tend to hold a much smaller buffer of liquid wealth to insulate them from the shock than do homeowners.

As discussed earlier, we expect that bank health will matter most for the consumption of constrained households. Renters’ consumption is likely to be more sensitive to bank health as well. On average, renters have lower income than homeowners and thus are more likely to live paycheck to paycheck and potentially need credit to smooth their consumption. Renters also tend to be less established financially and in their chosen profession, suggesting they may experience more variable income and more uneven expenses. Indeed, renters have a higher coefficient of variation (the ratio of the standard deviation to the mean) of income (0.64) than do owners (0.57)—see Table 2. Finally, the much lower level of liquid wealth held by renters (on average) makes them more susceptible to a reduction in credit availability due to a deterioration in bank health.

5.1 Initial Estimates

Our initial estimates of equation (1) show the expected negative relationship between bank health and the real food consumption of owners and renters (see Table 3, column 2), although in this specification the estimates are not statistically significant. (Recall that we report absolute effects.) As anticipated, the estimated coefficient for renters is somewhat larger (in absolute value) than that for homeowners, although the difference is not significant due to the imprecision with which the coefficients are estimated. In addition, including bank health in the analysis does not noticeably alter the other estimated relationships with consumption, such as income or homeowners’ propensity to consume out of their housing equity (see column 1).²³

The coefficient estimates for the non-bank-health variables are also what we would expect: Renters’ elasticity of consumption with respect to income is somewhat higher than homeowners’ elasticity, with the difference being statistically significant at all conventional levels. Renters’ greater sensitivity to income is

²³ The point estimates for the insignificant effects of local HPG and BEG diminish somewhat.

consistent with their being more likely to have spending dictated by current income. In addition, homeowners' marginal propensity to consume out of housing equity is about 10.8 cents per additional dollar,²⁴ which is on the higher side of existing estimates. Local HPG has a statistically insignificant effect on owners' and renters' consumption, and the effect of local employment growth is negative, likely due to collinearities with other controls. For presentation purposes, we show estimates for only the most relevant regressors in Table 3, while Table A1 in the Appendix shows our full set of results. The estimated effects for our other demographic control variables are also consistent with what we would expect based on life-cycle consumption theory.

When we focus on the sample for which we have financial wealth data (column 3), the bank health effects for owners and renters are much larger (in absolute value) and are highly significant, whereas most of the other estimated effects do not change much. The estimated bank health effect for renters increases the most (in absolute value) and is nearly double the effect for owners, although the difference is not significant at conventional levels. This finding is consistent with the notion that renters are more sensitive to credit availability than are homeowners. Controlling for households' financial wealth holdings directly (column 4) has virtually no impact on the estimated bank health effects, even though financial wealth itself has a positive and significant effect on food consumption for renters and owners. Accordingly, the sharp increase in the bank health effect between columns 2 and 4 is due to the difference in the coverage of the two samples, with the financial wealth subsample being much smaller (about half the size) and more concentrated in recent years. Overall, our initial results suggest that, especially for our financial wealth sample, bank health has an independent impact on household consumption beyond the effects of household wealth and income—an impact that appears to be greater for renters than for owners.²⁵

5.2 Bank Health and Constrained Households

5.2.1 Constrained: Negative Income Deviation Definitions

When we examine how our estimated bank health effects depend on whether households are likely constrained and in need of (bank-financed) credit, we find that bank health matters more for constrained households' food spending than unconstrained households' food spending.

Using any negative income deviation (shortfall) ($dev_{i,j} < 0$) to define constrained households (Table 4,

²⁴ Converting from the estimated (re-scaled) housing-equity effect in Table 3 (0.0118) to an MPC requires dividing by the ratio of average real food consumption to after-tax income over the sample (0.109).

²⁵ The implied MPC out of housing equity is 7.6 cents when we control for financial wealth (column 4)—an effect that is more consistent with estimates in the existing literature.

columns 1 and 2)²⁶, we find that a deterioration in bank health has a negative impact on consumption for both constrained owners and constrained renters—impacts that are significant at the 10 percent and 5 percent levels, respectively. Worse bank health has less of an effect (in absolute value) on unconstrained owners’ and renters’ consumption, with the difference between unconstrained and constrained owners more precisely estimated than the difference between unconstrained and constrained renters (p-value 0.03 versus 0.18). Moreover, the bank health effects are larger (in absolute value) for constrained renters than for constrained owners, which is consistent with the idea that renters are less secure financially. However, this difference is not statistically significant. Overall, households subjected to adverse income shocks appear more dependent on local banks’ ability and willingness to lend in order to smooth consumption over time. In addition, it appears that constrained homeowners rely on their home equity to smooth their consumption, as well as other non-housing-related bank financing, because the estimated effect of home equity on consumption for constrained owners is double that for unconstrained owners.

Our results are similar when we define constrained households based on a negative income deviation of 5 percent or more (Table 4, columns 3 and 4).²⁷ In particular, the estimated effects of bank health on consumption are now negative for unconstrained as well as constrained owners and renters, with the effects greater (in absolute value) for constrained households. This finding is also consistent with the idea that consumption for households most likely needing to borrow is more sensitive to the amount of available credit from banks, all else being equal. However, there is now little difference between the estimated effects for constrained owners and constrained renters.

Interpretation

The point estimates for the bank health effects for constrained households in column 2 of Table 4 imply that, all else being equal, going from the 25th to 75th percentile of the (overall) RENPL10 distribution (1.3 percentage points) leads to roughly 2.2 percent lower annual real food consumption for renters and 1.2 percent lower consumption for owners, relative to each group’s average expenditures.²⁸ These effects amount to about \$94 less of real food consumption for renters and \$70 less of real food consumption for

²⁶ In Table 4 and elsewhere, we present estimates for constrained and unconstrained households in separate columns for ease of comparison. However, the effects for constrained and unconstrained households are estimated in the same regression so that the relevant estimated fixed effects, macro controls, and demographic controls are the same as a household switches between being constrained or unconstrained from year to year.

²⁷ The sample size is smaller for these estimates because households with a current income shortfall between 0 percent and 5 percent are excluded from the analysis.

²⁸ Using the overall RENPL10 interquartile range for comparing the effects of a change in bank health on renters and owners is appropriate for comparing the effect of the same-sized shock. However, this likely understates the difference in the impact of bank health on constrained owners and renters since the interquartile range for renters is larger.

homeowners. If, instead, we use our estimates where we define constrained households based on a 5 percent or greater negative income deviation (column 4), the implied consumption effects for constrained owners are about twice as large and the implied effects for constrained renters are about 20 percent greater. The implied effects for unconstrained renters or owners (using the estimates in either column 1 or 3) are much smaller and, statistically, are not distinguishable from there being no effect.

In absolute terms, bank health has a relatively small dollar impact on constrained households' food consumption. Still, the effects are not trivial. Based on the coefficients in column 2, the estimated dollar effects are equivalent to the impact on food consumption of a roughly 14 percent decline in income, relative to the mean, for renters and a roughly 7 percent decline in income, relative to the mean, for owners. In addition, most food expenditures tend to be relatively less discretionary than other spending categories (although the additional cost of eating out is discretionary), so households are more likely to adjust their spending elsewhere in response to a decrease in credit availability (or income). Therefore, these estimates likely represent a lower bound on the overall bank health effects. The bank health estimates also capture effects that are averaged across households, including some that likely need to borrow to maintain their desired consumption path and some that do not, given the differences across households in their liquid asset holdings.

5.2.2 Constrained: Negative Income Deviation and Age

When we augment our negative income deviation definitions for constrained households with information about households' age, we find results similar to when we focus on households' negative income deviations alone. The specifications in columns 5 through 8 of Table 4 add young households to our definitions of constrained, regardless of whether these households experience an income shortfall. In particular, we label a single-earner household constrained if the worker is younger than 25 years old; a dual-earner household is constrained if both earners are younger than 25. This approach incorporates the possibility that young households may wish to borrow to smooth their intertemporal consumption path due to having higher expected future income relative to their current income, regardless of whether their current income is below its predicted level.

Once again, we observe that bank health impacts constrained households' food consumption much more than unconstrained households' food consumption, and the effect for constrained renters is somewhat greater than for constrained owners, although the difference is not statistically significant. The main difference when we add young households to our definition of constrained is that we gain some power. Whereas only

the difference in the bank health effects on constrained owners versus unconstrained owners was precisely estimated for the constrained definition that did not include young households, we can now observe that the difference between constrained renters and unconstrained renters is more significant when young households are included (p-values 0.06 and 0.101 between columns 7 and 8, and 5 and 6, respectively).

Overall, the results in Table 4 are consistent with bank health mattering the most for the consumption of households that have the greatest borrowing needs. Even constrained homeowners' spending is sensitive to changes in local-bank health, most likely because banks play a central role in homeowners accessing equity in their homes through home equity loans, in addition to providing other types of consumer loans. In times of financial stress, banks may be more willing to lend against collateral than to provide unsecured loans. Still, our results suggest that homeowners rely on banks for other types of loans besides home equity, or that banks cut back even on collateralized loans in periods of distress.

We repeat the analysis in Table 4 using the sample of households for which we have data on financial wealth (see Table A2 in the Appendix). Generally, the estimated bank health effects are greater (in absolute value) than those in Table 4, with all estimated coefficients negative and with some of the effects for unconstrained owners being significant, suggesting that households' consumption is sensitive to periods of weakened bank health, even in the absence of negative income shocks. However, bank health effects still tend to be greater (in absolute value) for constrained households than for unconstrained households and for renters relative to owners. These results confirm that our baseline bank health findings for constrained versus unconstrained households are not simply serving as a proxy for differences in households' financial resources.

5.2.3 Negative Income Deviation and Liquid Wealth Definitions

When we further refine our definition of constrained based on households' liquid wealth (LW) holdings, we find that, as expected, the consumption of constrained households with low LW is most sensitive to fluctuations in bank health (see Table 5). This result is consistent with low LW households having less savings easily available to serve as an alternative to bank financing during periods of income shortfalls.

More specifically, the specifications in Table 5 subdivide both the constrained and unconstrained household groups (based on our 5 percent negative deviation definition) between households with high (above median) LW and those with low (below median) LW.²⁹ Note that the sample size for this analysis is smaller than for our results in Table 4 due to the more limited number of PSID waves for which we have financial

²⁹ Liquid wealth includes cash holdings (checking and savings accounts), bond holdings, and stocks.

assets information. Further dividing the sample based on high and low LW leaves some groups—especially renters—with a fairly small number of observations in some cells, which will affect the power of our estimates.

Despite the smaller sample, when we control for LW, the results are consistent with what we would expect and with our earlier findings. Bank health has a negative estimated coefficient for each category of households, with larger (in absolute value) effects for constrained compared with unconstrained and for low LW compared with high LW households. The consumption of both constrained owners and constrained renters with low LW is particularly sensitive to changes in bank health, suggesting that given their limited liquid resources, both types of households depend on bank credit to smooth their expenditures over time.³⁰ The consumption of unconstrained households with low LW is also somewhat sensitive to bank health, all else being equal. Though only the effects for owners are precisely estimated, the magnitude of the bank health effects for both owners and renters in this category is reasonably large, suggesting that liquid wealth matters for smoothing consumption, perhaps due in part to expenditure shocks. That is, households with less liquid wealth to fall back on in periods of weak bank health have consumption that is more sensitive to bank health regardless of whether they are likely constrained due to an income shortfall.

The consumption of constrained households with a larger LW buffer is also sensitive to bank health (column 1). These households are sensitive to bank credit availability, either because they do not have sufficient liquid resources to cover the shortfall or they behave like “wealthy hand-to-mouth” consumers (see Kaplan, Violante, and Weidner 2014), taking on debt in order to maintain a sufficient buffer of liquid wealth for their possible future needs. Overall, the results in Table 5 are consistent with bank health impacting household consumption beyond the standard income and wealth channels, especially for households that are most likely to need credit but do not have a large liquid asset cushion to help protect them from an adverse income shock.

6 Robustness Tests

6.1 Alternative Measures of Bank Health

Supervisory Data

³⁰ However, unlike most of the previous results, the point estimates of the bank health effects are slightly larger (in absolute value) for low LW owners than for low LW renters, whether constrained or unconstrained. In addition, the bank health effect for constrained owners with low LW is statistically different from that for constrained owners with high LW. The same is not true for constrained renters.

When we use an alternative measure of bank health based on supervisory (CAMELS) data rather than bank balance sheet data, the results are similar to our earlier results. Indeed, the estimates are consistent with what we would expect, given the borrowing needs of renters versus owners and those of constrained versus unconstrained households (see Table 6). In particular, the sensitivity of renters' food consumption to changes in bank health, as captured by the supervisory data, is more than double (in absolute value) the sensitivity of owners' consumption—a difference that is statistically significant at the 10 percent level. While the bank health effects are stronger for constrained owners compared with unconstrained owners, as they are in Table 4, the effects for constrained versus unconstrained renters are of similar size. This divergence in results for renters relative to those in Table 4 is likely due to the somewhat different sample horizon for the CAMELS data.

Indeed, when we switch to the more comparable sample for which we have financial wealth data (columns 4 through 6), the results are in line with what we observed earlier. The sensitivity of food consumption to bank health is greater for constrained renters than unconstrained renters as well as for constrained versus unconstrained owners, and both differences are statistically significant. The bank health effect for constrained renters is also greater (more negative) than the effect for constrained owners, consistent with renters having fewer resources to fall back on when bank credit is less readily available, although this difference is not precisely estimated.

Multi-locational Bank Health

We observe a pattern of results similar to our main findings in Table 4 when we use our ML BH measure, which is arguably more exogenous with respect to local economic conditions. Recall that by definition these multi-locational banks are relevant (in terms of deposits) for households in the local area, but their presence is limited in the context of the banks' broader footprint. The results in Table 7 show that renters' food consumption is more than twice as sensitive to changes in ML BH compared with owners' food consumption (column 1). We also observe the familiar pattern in which constrained households' consumption depends much more on fluctuations in bank health than the consumption of unconstrained households—for both renters and owners (columns 2 and 3). The estimated ML BH effect for constrained renters is also greater in absolute value than for constrained owners. The results are qualitatively similar if we restrict our analysis to the sample for which we have financial wealth data (columns 4 through 6); therefore, differences in ML BH are not serving as a proxy for differences in households' financial assets. Our findings are also similar if we use a multi-locational measure of bank health based on the balance sheet of a branch's parent bank rather

than the balance sheet of its holding company (not shown).

Overall, the results in Tables 6 and 7 suggest that our findings are not driven by our approach to measuring local-bank health. That is, our bank health results do not appear to also be capturing differences in local economic conditions that may impact local-bank health.

6.2 Alternative Measures of Consumption

Our results exhibit similar patterns if we replace food consumption with alternative, broader measures of household expenditures. Columns 1 through 3 in Table 8 report results using an imputed (nondurables) measure of consumption based on the approach in Blundell, Pistaferri, and Preston (2006). This approach uses the estimated relationship between nondurables consumption, food consumption, and observable household characteristics in the CEX to infer nondurables consumption in the PSID. Qualitatively, the effects are similar to our baseline findings. Bank health continues to matter most for constrained households' consumption, as we would expect, and the effect is slightly greater for constrained renters than for constrained homeowners. In fact, the bank health effects for constrained households (column 3) are the only ones that are statistically significant. In comparison, the bank health effects for unconstrained renters and unconstrained owners are much smaller. Moreover, the bank health effects for the full sample have the predicted negative signs.

Columns 4 through 6 in Table 8 show the results when we use the additional consumption data that are available in the PSID starting in 1999 (see Appendix Section A.1.1). Whereas this measure of expenditures is broader than food consumption, it is available over a much shorter time horizon—a period that excludes the banking crisis of the late 1980s and early 1990s. Therefore, we likely have less power for identifying the bank health effects. That said, the bank health effects for owners are significant, but the effect for unconstrained owners is larger in magnitude than for constrained owners, although this difference is not statistically significant. Among renters, we continue to find a large bank health effect for those who are constrained and likely need to borrow, but it is imprecisely estimated. While it is difficult to interpret and compare these results due to the shorter sample and exclusion of the banking crisis, we continue to find evidence that household spending is sensitive to bank health.

7 Conclusion

The financial crisis and related contraction in bank credit associated with the Great Recession highlighted the need for a better understanding of the links between financial markets and the macroeconomy.

From a financial stability perspective, it is important for policymakers to understand the linkages between the financial sector and real activity. While a number of papers explore this relationship, this study focuses on the more general effects of reduced credit availability operating through consumer loans rather than primarily emphasizing the effects on homeowners operating through the loss of home equity. We examine the relationship between the health of local financial institutions (banks) and household-level consumer spending using an innovative approach to identifying households that most likely are constrained and are most in need of accessing bank credit to smooth their consumption. Even though there has been much consolidation in the banking sector over time, and primary mortgages today are underwritten by national lenders, banks with local branches still play an important role in providing home equity and other consumer loans.

Using data on the health of banks with a local presence and the spending behavior of households in those locations, this study examines the relationship between local-bank health and consumption, paying particular attention to differences in the spending responses to fluctuations in bank health of homeowners versus renters and of constrained versus unconstrained households. The results show that while a deterioration in bank health results in lower real (food) consumption, the effects tend to be strongest for constrained households—especially constrained renters, as well as owners and renters with limited liquid wealth—which have limited resources to fall back on in response to an income shock or contraction in bank credit. The study further shows that these results are robust to using supervisory data instead of balance sheet data to determine bank health and to measure the health of multi-locational banks whose limited presence in a given location makes their health arguably exogenous relative to local economic conditions.

References

- Abdallah, Chadi, and William Lastrapes. 2012. "Home Equity Lending and Retail Spending: Evidence from a Natural Experiment in Texas." *American Economic Journal: Macroeconomics* 4(4): 94–125. Available at <http://pubs.aeaweb.org/doi/pdfplus/10.1257/mac.4.4.94>.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2018. "Do Banks Pass through Credit Expansions to Consumers Who Want to borrow?" *Quarterly Journal of Economics* 133(1): 129–190. Available at <https://doi.org/10.1093/qje/qjx027>.
- Agarwal, Sumit, and Wenlan Qian. 2017. "Access to Home Equity and Consumption: Evidence from a Policy Experiment." *Review of Economics and Statistics* 99(1): 40–52. Available at http://www.mitpressjournals.org/doi/abs/10.1162/REST_a_00606.
- Albuquerque, Bruno, and Georgi Krustev. 2018. "Debt Overhang and Deleveraging in the US Household Sector: Gauging the Impact on Consumption." *Review of Income and Wealth* 64(2): 459–481. Available at <https://onlinelibrary.wiley.com/doi/abs/10.1111/roiw.12265>.
- Argyle, Bronson, Taylor Nadauld, and Christopher Palmer. 2017. "Real Effects of Search Frictions in Consumer Credit Markets." Available at <https://www.fdic.gov/news/conferences/consumersymposium/2017/documents/argyle-paper.pdf>.
- Bartik, Timothy. 1991. *Who Benefits from State and Local Development Policies?* Kalamazzo, MI: W.E. Upjohn Institute for Employment Research.
- Bassett, William, Mary Beth Chosak, John Driscoll, and Egon Zakrajsek. 2014. "Changes in Bank Lending Standards and the Macroeconomy." *Journal of Monetary Economics* 62: 23–40. Available at <https://www.sciencedirect.com/science/article/pii/S0304393213001645>.
- Berrospide, José M., Lamont K. Black, and William R. Keeton. 2016. "The Cross - Market Spillover of Economic Shocks through Multimarket Banks." *Journal of Money, Credit, and Banking* 48(5): 957–988. Available at <https://onlinelibrary.wiley.com/doi/abs/10.1111/jmcb.12323>.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2006. "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX." Available at <https://web.stanford.edu/~pista/impute.pdf>.
- Board of Governors of the Federal Reserve System. 2018. "Report on the Economic Well-Being of U.S. Households in 2017." Available at <https://www.federalreserve.gov/publications/files/2017-report-economic-well-being-us-households-201805.pdf>.
- Campbell, John, and Joao Cocco. 2007. "How Do House Prices Affect Consumption? Evidence from Micro Data." *Journal of Monetary Economics* 54(3): 591–621. Available at <https://doi.org/10.1016/j.jmoneco.2005.10.016>.
- Charles, Kerwin, and Erik Hurst. 2003. "The Correlation of Wealth across Generations." *Journal of Political Economy* 111(6): 1155–1182. Available at <https://www.journals.uchicago.edu/doi/10.1086/378526>.
- Chava, Sudheer, and Amiyatosh Purnanandam. 2011. "The Effect of Banking Crisis on Bank-Dependent Borrowers." *Journal of Financial Economics* 99(1): 116–135. Available at <http://www.sciencedirect.com/science/article/pii/S0304405X10001820>.

- Cooper, Daniel. 2013. "House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral." *Review of Economics and Statistics* 95(4): 1183–1197. Available at http://www.mitpressjournals.org/doi/abs/10.1162/REST_a_00323.
- Cooper, Daniel, and Karen Dynan. 2016. "Wealth Effects and Macroeconomic Dynamics." *Journal of Economic Surveys* 30(1): 34–55. Available at <https://EconPapers.repec.org/RePEc:bla:jecsur:v:30:y:2016:i:1:p:34-55>.
- Cortés, Kristle, and Phillip Strahan. 2015. "Tracing out Capital Flows: How Financial Integrated Banks Respond to Natural Disasters." Federal Reserve Bank of Cleveland Working Paper 14–21. Available at <https://www.clevelandfed.org/newsroom-and-events/publications/working-papers/2014-working-papers/wp-1412-tracing-out-capital-flows-how-financially-integrated-banks-respond-to-natural-disasters.aspx>.
- Damar, Evren H., Reint Gropp, and Adi Mordel. 2014. "Banks' Financial Distress, Lending Supply and Consumption Expenditure." Bank of Canada Staff Working Paper Series 2014-17. Available at <http://www.bankofcanada.ca/2014/02/working-paper-2014-7/>.
- Dynan, Karen. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity*, Spring 2012: 299–362. Available at www.jstor.org/stable/23287219.
- Greenstone, Michael, Alexandre Mas, and Hoai-Luu Q. Nguyen. 2014. "Do Credit Market Shocks Affect the Real Economy? Quasi-Experimental Evidence from the Great Recession and 'Normal' Economic Times." NBER Working Papers 20704. Available at <https://ideas.repec.org/p/nbr/nberwo/20704.html>.
- Hurst, Erik, Annamaria Lusardi, Arthur Kennickell, and Francisco Torralba. 2010. "The Importance of Business Owners in Assessing the Size of Precautionary Savings." *Review of Economics and Statistics* 92(1): 61–69. Available at <https://www.mitpressjournals.org/doi/10.1162/rest.2009.11500>.
- Hurst, Erik, and Frank Stafford. 2004. "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption." *Journal of Money, Credit, and Banking* 36(6): 985–1014. Available at http://econpapers.repec.org/article/mcbjmoncb/v_3a36_3ay_3a2004_3ai_3a6_3ap_3a985-1014.htm.
- Jensen, Thais, and Niels Johannesen. 2017. "The Consumption Effects of the 2007 – 2008 Financial Crisis: Evidence from Households in Denmark." *American Economic Review* 107(11): 3386–3414. Available at <https://www.aeaweb.org/articles?id=10.1257/aer.20151497>.
- Johnson, David, Jonathan Parker, and Nicholas Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96(5): 1589–1610. Available at <https://www.aeaweb.org/articles?id=10.1257/aer.96.5.1589>.
- Johnson, Kathleen, and Geng Li. 2010. "The Debt-Payment-to-Income Ratio as an Indicator of Borrowing Constraints: Evidence from Two Household Surveys." *Journal of Money, Credit and Banking* 42(7): 1373–1390. Available at <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1538-4616.2010.00345.x>.
- Kaplan, Greg, Giovanni Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Paper on Economic Activity*, Spring 2014: 77–153. Available at https://www.brookings.edu/wp-content/uploads/2016/07/2014a_Kaplan.pdf.
- Kazarosian, Mark. 1997. "Precautionary Savings – A Panel Study." *Review of Economics and Statistics* 79(2): 241–247. Available at <https://www.mitpressjournals.org/doi/10.1162/003465397556593>.

- Lehnert, Andreas. 2004. "Housing, Consumption, and Credit Constraints." Board of Governors of the Federal Reserve System Finance and Economic Discussion Series No. 2004-63. Available at <http://www.federalreserve.gov/pubs/feds/2004/200463/200463pap.pdf>.
- Mian, Atif, Kamalesh Rao, and Amri Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *The Quarterly Journal of Economics* 128(4): 1687–1726. Available at <https://academic.oup.com/qje/article/128/4/1687/1849337>.
- Muellbauer, John. 2008. "Housing, Credit and Consumer Expenditure." CEPR Discussion Paper No. 6782. Available at <https://econpapers.repec.org/paper/cprceprdp/6782.htm>.
- Nguyen, Hoai-Luu Q. Forthcoming. "Are Credit Markets Still Local? Evidence from Bank Branch Closings." *American Economic Journal: Applied Economics* Available at <https://www.aeaweb.org/articles?id=10.1257/app.20170543&&from=f>.
- Zeldes, Stephen. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2): 305–346. Available at <https://www.journals.uchicago.edu/doi/10.1086/261605/>.

Table 1: Bank Health and Annual Real Per Capita Nonmortgage Loan Growth

	Excluding Home Equity		Including Home Equity	
	(1)	(2)	(3)	(4)
RENPL10/100	−18.939*** (4.359)		−43.389*** (5.913)	
RENPL10/100 Multi-BHC (10%)		−13.673* (7.100)		−46.685*** (8.966)
House Price Growth (%)	0.091*** (0.008)	0.096*** (0.008)	0.176*** (0.012)	0.180*** (0.013)
Bartik Empl. Growth (%)	0.497*** (0.134)	0.430*** (0.140)	0.517*** (0.154)	0.495*** (0.166)
Obs.	10616	9606	10616	9606
Adj. R-squared	0.504	0.515	0.510	0.517

Sources: Authors' calculations using the NY Fed Consumer Credit Panel/Equifax (CCP), Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Annual data at the Metropolitan Statistical Area (MSA) level. Sample restricted to MSAs found in the PSID data. Sample that includes home equity is based on the subset of the CCP individuals with mortgage loan data (see main text for more details). Additional controls include location (MSA) fixed effects and time fixed effects. The dependent variable in the first two columns is the percent change in real per capita nonmortgage loans. The dependent variable in the last two columns is the percent change in real per capita nonmortgage loans plus home equity loans. RENPL10 is the baseline balance sheet (nonperforming real estate loan ratio) measure of bank health at the MSA level (see Section 2.2). Multi-locational banks are bank holding companies with less than 5 percent of their overall deposits held in that location, but where those deposits account for at least a 10 percent share of all local deposits. Robust standard errors are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 2: Summary Statistics for Baseline Sample

Variable	Renters						Owners					
	Mean	Median	Std. Dev	Min.	Max.	N	Mean	Median	Std. Dev	Min.	Max.	N
Food Consumption	4,272	3,826	3,081	11	55,039	39,888	6,136	5,749	3,458	28	40,608	53,188
Broader PSID Consumption Measure	15,695	13,174	12,305	471	102,777	16,238	26,156	22,927	15,608	2,004	120,017	23,286
Imputed Nondurable Consumption	10,923	10,113	6,549	115	78,574	33,431	14,459	13,779	7,050	175	59,974	49,661
Income	30,200	26,567	19,390	667	138,188	39,888	56,451	50,778	31,953	2,509	300,554	53,188
RENPL10/100	.014	.0092	.014	.0013	.099	39,888	.014	.0085	.013	.0013	.099	53,188
RENPL10/100: Multi-BHC (10%)	.015	.0085	.015	0	.17	19,839	.014	.0076	.015	0	.17	28,923
CAMELS	2	1.9	.43	1	4.4	33,294	1.9	1.9	.41	1	4.4	45,510
Home Equity/Income	0	0	0	0	0	39,888	1.3	.86	1.5	-1.3	20	53,188
Fin. Wealth/Income	.12	.043	.71	-4.4	7.5	19,096	1.1	.39	2.2	-1.9	27	25,991
House Price Growth (%)	1.2	1.1	6.6	-30	34	39,888	1.1	1.1	6.5	-30	34	53,188
Bartik Empl. Growth (%)	1.7	2.2	1.8	-5.9	5.1	39,888	1.6	2	1.9	-5.9	5.1	53,188
Age	37	34	11	17	88	39,888	45	44	11	18	92	53,188
Family Size	2.7	2	1.6	1	14	39,888	3.1	3	1.4	1	14	53,188
Single	.59	1	.49	0	1	39,888	.23	0	.42	0	1	53,188
Single Earner	.13	0	.33	0	1	39,888	.23	0	.42	0	1	53,188
Dual Earner	.28	0	.45	0	1	39,888	.54	1	.5	0	1	53,188
Liquid Wealth/Income	.11	.0036	.37	0	6.2	19,096	.46	.086	1.1	0	27	25,991
Percent of Const. HHs (Neg. Dev.)	51	51	4.4	40	57	26,353	48	50	5.7	38	55	43,841
Percent of Const. HHs ($\geq 5\%$ Neg. Dev.)	45	46	4.4	34	52	23,567	41	41	5.1	31	46	38,229

Sources: Authors' calculations using PSID, FDIC Summary of Deposits, Federal Reserve System, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: The broader PSID consumption measure is based on additional expenditure data included in the PSID starting in 1999. Imputed nondurable consumption is based on the approach in Blundell, Pistaferri, and Preston (2006). Income is after-tax real household income. RENPL10 is the baseline balance sheet (nonperforming real estate loan ratio) measure of bank health (see Section 2.2). CAMELS is the deposit-weighted average of bank supervisors' confidential bank ratings (integer scale of 1 to 5). Multi-locational banks are bank holding companies with less than 5 percent of their overall deposits held in that location, but where those deposits account for at least a 10 percent share of all local deposits. Home Equity/Income is the ratio of housing wealth (home value less outstanding mortgage debt) to lagged income. Fin. Wealth/Income is the ratio of financial wealth (bond, stock, cash, and retirement account holdings along with the net value of any business, vehicles, or non-primary-residence real estate properties owned) to lagged income. Renter Dummy is an indicator variable for whether the household rents or owns a home. Single Earner and Dual Earner are indicator variables for whether the household head is cohabitating and living in a single-earner household, living in a dual-earner household, or is single (excluded category). Liquid Wealth/Income is the ratio of liquid wealth (cash holdings (checking and savings accounts), stock holdings, and bond holdings) to lagged income. Percent of constrained households is the share of households in a given year that are constrained relative to the number of households with nonmissing data for the constraint indicator (see main text for details on how constrained households are defined).

Table 3: Bank Health and Consumption: Initial Results

	(1) Food	(2) Food	(3) Food	(4) Food
Log Income: Owner	0.191*** (0.013)	0.191*** (0.013)	0.221*** (0.013)	0.223*** (0.013)
Log Income: Renter	0.242*** (0.013)	0.242*** (0.013)	0.296*** (0.015)	0.297*** (0.015)
RENPL10/100: Owner		-0.632 (0.424)	-1.534*** (0.486)	-1.550*** (0.486)
RENPL10/100: Renter		-0.866 (0.638)	-2.946*** (1.016)	-2.945*** (1.015)
Home Equity/(Inc.*100): Owner	1.176*** (0.387)	1.166*** (0.388)	1.003** (0.425)	0.832** (0.421)
Fin. Wealth/(Inc.*100): Owner				0.979*** (0.189)
Fin. Wealth/(Inc.*100): Renter				2.370*** (0.855)
Renter Dummy	-0.759*** (0.181)	-0.759*** (0.180)	-0.758*** (0.231)	-0.734*** (0.229)
House Price Growth: Owner	0.064 (0.062)	0.032 (0.060)	0.042 (0.061)	0.042 (0.062)
House Price Growth: Renter	0.059 (0.101)	0.018 (0.097)	-0.056 (0.120)	-0.057 (0.120)
Bartik Empl. Growth: Owner	0.013 (0.584)	-0.049 (0.586)	-0.051 (0.664)	0.001 (0.667)
Bartik Empl. Growth: Renter	-1.234 (1.069)	-1.252 (1.072)	-2.288 (1.605)	-2.187 (1.613)
Owner Obs.	53188	53188	25991	25991
Renter Obs.	39888	39888	19096	19096
Total Obs.	93076	93076	45087	45087
Adj. R-squared	0.593	0.593	0.573	0.574

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is log real food consumption. Income is real after-tax household income. RENPL10 is the baseline balance sheet (nonperforming real estate loan ratio) measure of bank health (see Section 2.2). Home Equity/Inc. is the ratio of household housing wealth (home value less outstanding mortgage debt) to lagged income. Fin. Wealth/Inc. is the ratio of household financial wealth (bond, stock, cash, and retirement account holdings along with the net value of any business, vehicles, or non-primary-residence real estate properties owned) to lagged income. Renter Dummy is an indicator variable for whether the household rents or owns a home. Specifications include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): age, age squared, family size, family size squared, and indicators for single-earner and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 4: Bank Health and Consumption: Controlling for Constrained Households

	Negative Income Deviation		≥5% Neg. Dev.		Neg. Dev. or < 25		≥5% Neg. Dev. or < 25	
	(1) Unconst.	(2) Const.	(3) Unconst.	(4) Const.	(5) Unconst.	(6) Const.	(7) Unconst.	(8) Const.
Log Income: Owner	0.167*** (0.015)	0.166*** (0.015)	0.163*** (0.016)	0.163*** (0.017)	0.171*** (0.015)	0.171*** (0.015)	0.169*** (0.016)	0.169*** (0.017)
Log Income: Renter	0.178*** (0.016)	0.173*** (0.016)	0.175*** (0.017)	0.170*** (0.018)	0.183*** (0.016)	0.179*** (0.016)	0.182*** (0.017)	0.178*** (0.018)
RENPL10/100: Owner	0.089 (0.484)	-0.908* (0.495)	-0.193 (0.499)	-1.877*** (0.537)	0.117 (0.488)	-0.921* (0.503)	-0.160 (0.502)	-1.891*** (0.546)
RENPL10/100: Renter	-0.637 (0.719)	-1.669** (0.753)	-0.778 (0.746)	-1.986** (0.845)	-0.482 (0.725)	-1.759** (0.743)	-0.571 (0.751)	-2.142** (0.831)
Home Equity/(Inc.*100): Owner	0.461 (0.522)	0.939** (0.404)	0.413 (0.529)	0.832* (0.476)	0.370 (0.548)	0.986** (0.396)	0.320 (0.552)	0.896* (0.460)
Owner Obs.	43841		38229		43841		38229	
Renter Obs.	26353		23567		26353		23567	
Total Obs.	70194		61796		70194		61796	
Adj. R-squared	0.613		0.607		0.613		0.607	

Source: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is log real food consumption. Constrained households in column 2 are those with a negative income deviation, and constrained households in column 4 are those with a 5 percent or greater negative income deviation, as discussed in the main text. In columns 6 and 8, constrained households are those identified by the approach in columns 2 and 4, respectively, as well as any households in which all earners are younger than 25 and have valid age-earnings profiles. Unconstrained households in all specifications are those households without a negative income deviation. See Table 2 for additional variable definitions. All specifications also include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house-price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 5: Bank Health and Consumption: Controlling for Constrained Households and Liquid Wealth Holdings

	$\geq 5\%$ Neg. Dev.			
	(1)	(2)	(3)	(4)
	Const., High LW	Const., Low LW	Unconst., High LW	Unconst., Low LW
Log Income: Owner	0.189*** (0.016)	0.189*** (0.016)	0.188*** (0.015)	0.189*** (0.015)
Log Income: Renter	0.231*** (0.023)	0.228*** (0.023)	0.234*** (0.022)	0.235*** (0.022)
RENPL10/100: Owner	-2.557*** (0.757)	-4.441*** (1.085)	-1.027* (0.579)	-2.462*** (0.775)
RENPL10/100: Renter	-3.580** (1.621)	-4.002*** (1.422)	-1.526 (1.353)	-1.846 (1.253)
Home Equity/(Inc.*100): Owner	0.019 (0.605)	1.122 (0.890)	1.493** (0.662)	0.181 (1.059)
Fin. Wealth/(Inc.*100): Owner	1.102*** (0.339)	2.101** (0.903)	0.707*** (0.269)	0.821 (0.825)
Fin. Wealth/(Inc.*100): Renter	2.641 (1.705)	10.157*** (2.935)	0.045 (1.396)	-0.683 (2.839)
Owner Obs.		19085		
Renter Obs.		10707		
Total Obs.		29792		
Adj. R-squared		0.548		

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is real log food consumption. Constrained households have a 5 percent or greater negative income deviation as discussed in the main text. Constrained and unconstrained households are further divided by their liquid wealth holdings: high (above median) or low (below median). Liquid wealth includes cash holdings (checking and savings accounts), stock holdings, and bond holdings. See Table 2 for additional variable definitions. Specifications also include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner, and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 6: Bank Health and Consumption: Supervisory Measure of Bank Health

	(1)	$\geq 5\%$ Neg. Dev.		(4)	$\geq 5\%$ Neg. Dev.	
	Food	Unconst.	Const.	Food	Unconst.	Const.
Log Income: Owner	0.181*** (0.015)	0.152*** (0.017)	0.154*** (0.017)	0.214*** (0.014)	0.177*** (0.016)	0.189*** (0.018)
Log Income: Renter	0.223*** (0.014)	0.166*** (0.019)	0.158*** (0.020)	0.283*** (0.015)	0.225*** (0.023)	0.232*** (0.025)
CAMELS: Owner	-0.032** (0.014)	-0.023 (0.021)	-0.054*** (0.017)	-0.044** (0.017)	-0.022 (0.021)	-0.099*** (0.025)
CAMELS: Renter	-0.079*** (0.027)	-0.082*** (0.025)	-0.073** (0.036)	-0.102*** (0.031)	-0.059 (0.040)	-0.145*** (0.038)
Home Equity/(Inc.*100): Owner	1.072** (0.475)	0.127 (0.642)	0.865 (0.548)	0.804* (0.433)	1.070 (0.667)	0.441 (0.534)
Fin. Wealth/(Inc.*100): Owner				0.835*** (0.217)	0.835*** (0.302)	1.078*** (0.344)
Fin. Wealth/(Inc.*100): Renter				2.114*** (0.916)	-0.948 (1.422)	5.729*** (1.602)
Owner Obs.	45510	33390		23660	17460	
Renter Obs.	33294	19684		17343	9591	
Total Obs.	78804	53074		41003	27051	
Adj. R-squared	0.595	0.611		0.577	0.551	

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, Federal Reserve System, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is real log food consumption. CAMELS is the deposit-weighted average of bank supervisors' confidential bank ratings (integer scale of 1 to 5) in the local area. Constrained households have a 5 percent or greater negative income deviation, and unconstrained households are those without a negative income deviation as discussed in the main text. See Table 2 for additional variable definitions. Specifications also include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house-price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 7: Bank Health and Consumption: Multi-Locational Measure of Bank Health

	≥5% Neg. Dev.		≥5% Neg. Dev.	
	(1)	(2)	(3)	(4)
	Food	Unconst.	Const.	Food
Log Income: Owner	0.185*** (0.017)	0.150*** (0.018)	0.151*** (0.018)	0.216*** (0.018)
Log Income: Renter	0.248*** (0.017)	0.188*** (0.029)	0.184*** (0.030)	0.297*** (0.021)
RENPL10/100: Owner	-0.942** (0.386)	-0.371 (0.592)	-2.234*** (0.662)	-0.749* (0.417)
RENPL10/100: Renter	-2.147*** (0.641)	-0.702 (0.820)	-3.510*** (0.873)	-3.307*** (1.056)
Home Equity/(Inc.*100): Owner	0.783 (0.516)	0.388 (0.629)	0.632 (0.805)	0.779 (0.551)
Fin. Wealth/(Inc.*100): Owner				0.696*** (0.253)
Fin. Wealth/(Inc.*100): Renter				1.404 (1.112)
Owner Obs.	28923	21133		16908
Renter Obs.	19839	11420		12457
Total Obs.	48762	32553		29365
Adj. R-squared	0.609	0.617		0.578
				12468
				6599
				19067
				0.551

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is real log food consumption. Multi-locational banks are bank holding companies with less than 5 percent of their overall deposits held in that location, but where those deposits account for at least a 10 percent share of all local deposits. See Table 2 for additional variable definitions. Specifications include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house-price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 8: Bank Health and Consumption: Alternative Measures of Household Expenditures

	(1)	≥5% Neg. Dev.		(4)	≥5% Neg. Dev.	
	Imputed	Unconst.	Const.	Broader PSID	Unconst.	Const.
Log Income: Owner	0.125*** (0.011)	0.104*** (0.012)	0.103*** (0.012)	0.204*** (0.013)	0.185*** (0.016)	0.180*** (0.017)
Log Income: Renter	0.176*** (0.010)	0.126*** (0.015)	0.123*** (0.015)	0.244*** (0.014)	0.244*** (0.022)	0.243*** (0.023)
RENPL10/100: Owner	-0.358 (0.306)	-0.155 (0.366)	-1.135*** (0.354)	-0.955*** (0.284)	-1.467*** (0.422)	-1.012*** (0.447)
RENPL10/100: Renter	-0.296 (0.505)	-0.393 (0.534)	-1.326** (0.625)	-0.396 (0.419)	-0.160 (0.640)	-1.261 (0.898)
Home Equity/(Inc.*100): Owner	0.437 (0.276)	-0.253 (0.431)	0.306 (0.409)	1.531*** (0.499)	1.369*** (0.649)	1.526*** (0.649)
Owner Obs.	49661	36244		23286		17445
Renter Obs.	33431	20543		16238		8861
Total Obs.	83092	56787		39524		26306
Adj. R-squared	0.648	0.661		0.642		0.612

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: In columns 1-3, the dependent variable is real log imputed nondurables expenditures based on the method in Blundell, Pistaferri, and Preston (2016). In columns 4-6, the dependent variable is real log household expenditures based on the additional expenditure data included in the PSID starting in 1999: food, healthcare, childcare, transportation, school, vehicle, and utilities. Constrained households have a 5 percent or greater negative income deviation, and unconstrained households are those without a negative income deviation, as discussed in the main text. See Table 2 for additional variable definitions. Specifications also include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner, and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Appendix

A.1 Additional Data Details

A.1.1 PSID Data

A few additional facts regarding the PSID data and the calculations we perform with the data are worth highlighting, given our analysis in the main text.

- Reported income data in each PSID wave cover household earnings for the previous calendar year. For example, income data reported in the 2011 wave are for calendar year 2010.
- Starting in 1999, the PSID expanded its recorded household expenditure categories to include health-care, childcare, transportation, school, vehicle, and utilities, in addition to food. We use the sum of these consumption categories when we test the robustness of our baseline estimates discussed in Section 6 that are based on food consumption alone.
- The PSID further expanded its recorded spending categories in 2005, but the time horizon of these data is too short to include them in our analysis. These additional categories include home maintenance (upkeep), clothing, recreation, home furnishings, and vacation expenditures.
- Where relevant, we use the quarterly personal consumption expenditures (PCE) deflator to convert nominal values to real values, such as for consumption and household income.
- We estimate each household's income tax burden using the National Bureau of Economic Research (NBER) tax simulation (TAXSIM) module.
- Housing wealth is defined as a household's self-reported home value less any outstanding (primary or secondary) mortgage debt. (Housing wealth is zero for renters). Financial wealth includes the value of a household's bond, stock, cash, and retirement account holdings, along with the net value of any businesses, vehicles, or non-primary-residence real estate properties owned. Financial wealth is net of any outstanding non-housing debt for the household.

A.1.2 Bank Sample Selection

We employ additional restrictions beyond what we discuss in the main text to generate an appropriate sample of banks from the Call Report data for our analysis.

- To capture active lenders, we eliminate non-lending institutions by dropping banks with an average share of total loans relative to total assets that is less than 10 percent.
- We exclude credit card banks—those banks that at any point in the sample period had a ratio of outstanding credit card loans to total loans greater than 50 percent.¹
- To address the issue of banks behaving abnormally immediately before voluntary liquidations or when they are newly established (de novo banks), we drop the final two quarters of observations of voluntarily liquidated banks and the first eight quarters in which a de novo bank files a Call Report.

A.1.3 Constructing Multi-Locational Bank Health

For a given location, we deem a bank multi-locational if less than 5 percent of the bank’s deposits come from that particular location. We then reconstruct our bank health measures for the set of multi-locational banks with branches in that location. However, we include in our regression sample only the locations in which multi-locational banks account for at least 10 percent of local deposits to ensure that these banks have a potentially meaningful effect on credit availability in that location. Note that each bank is determined to be, or not be, multi-locational on a location-by-location basis, because a given bank operating in multiple locations may have some locations that account for more than 5 percent of its deposits, while its branches in other locations may account for less than 5 percent. Also, because internal capital markets operate within bank holding companies, we construct the multi-locational bank health measures based on the affiliation of a given branch with the top holding company of the bank that owns the branch.²

A.1.4 Calculating House Price Growth

To calculate local house price growth, we use quarterly data from CoreLogic’s “single family combined index,” which includes distressed sales. We convert the index from nominal to real data using the PCE deflator and then define house price growth in a given location as the four-quarter percent change in the real house price index. Rural households or households in a metropolitan statistical area (MSA) with no

¹ Because credit card banks are disproportionately located in South Dakota and Delaware, they tend to distort bank health measures in those states. In addition, credit card banks tend to securitize credit card loans rather than relying on deposits to fund them. These banks do not necessarily lend locally and thus would distort our measures of local credit availability. We also remove a number of anomalous banks—those for which the “Bank Type Analysis Code” [Call Report variable rssid9425] does not equal zero.

² Because most of the substantial consolidation of the banking sector occurred after the Riegle-Neal Interstate Banking and Branching Efficiency Act of 1994, which allows banks to branch across state lines, the early part of our sample would have relatively few multi-locational banks if we instead defined them based on a branch’s affiliation with its bank rather than with the bank’s top holding company.

CoreLogic house price data are assigned state-level CoreLogic real house price growth rates, which we construct the same way. We assign house price growth rates to households with respect to their interview date (year and quarter), as discussed in Section 3.2 of the main text.

A.1.5 Calculating Local (Bartik) Employment Growth

We construct local Bartik employment growth (BEG) by first calculating each industry's share of employment in each MSA. We then weight the national four-quarter employment percent change for each industry by the lagged employment share of that industry in each MSA. Finally, we sum each of these weighted growth rates across industries to obtain the MSA-level BEG. We follow the same procedure to construct state-level BEG, which we assign to households in the non-MSA (rural) locations in each state.

The MSA and state-level industry employment data come from the Quarterly Census of Employment and Wages (QCEW), published by the Bureau of Labor Statistics (BLS).³ The QCEW data are available at the county level, and we sum employment in the counties within each MSA to obtain MSA-level measures of employment.⁴

In calculating BEG, we make some adjustments to the QCEW data, because industry classifications changed over time. In particular, the QCEW reports industry-based data according to the North American Industry Classification System (NAICS) from 2001 to the present and provides data based on the Standard Industrial Classification (SIC) system in earlier years. There are also SIC data from before 2001 that have been reconstructed using the NAICS industry classification approach. However, due to inaccuracies in the reconstructed data, we combine the original SIC data with the NAICS data to generate consistently measured industry shares by location over time.

Consistent industry shares require a level of industry employment that does not change based on the shift from the SIC classification system in 2000 to the NAICS system in 2001. Our approach for both the MSA-level and state-level data is as follows:

- We standardize industries at the county level between the two classification systems into 14 broad industries.⁵

³ See <https://www.bls.gov/cew/datatoc.htm> for more details on these data.

⁴ The advantage of using the county-level data is that it allows us to control for changes in the geographic boundaries of some MSAs as they occur. As necessary, we use available data to approximate county-level, state-level, and other relevant employment data that are not disclosed in a given quarter for confidentiality reasons.

⁵ The 14 industries are agriculture, forestry, fishing, and hunting; mining; utilities; construction; manufacturing; retail and wholesale trade; transportation and warehousing; information; finance, insurance, real estate, rental and leasing; professional and business services; educational services, healthcare, and social assistance; art, entertainment, recreation, accommodation, and food services; other services except government; government. Using a high level of industry aggregation minimizes any remaining inconsistencies between the two classification systems.

- We take local-level 2001:Q1 NAICS-based employment data by industry and extend it backward to the beginning of our sample using the relevant SIC-based industry growth rates.
- In cases where industry employment in 2001:Q1 is zero (or missing) in a location, we rescale the most recent non-zero SIC employment data prior to 2001 by the national ratio of SIC employment to reconstructed (national) NAICS employment in the industry at that time. We then grow this imputed value backward using the SIC growth rates.⁶

With this adjusted industry-level employment, we calculate industry shares by location and use them to weight national industry-level employment growth rates. However, to do so we also need to adjust the national employment growth rates between 2000 and 2001 to account for the shift in classification systems. In particular, we combine our reconstructed (national) NAICS data for 2000 with the actual NAICS data for 2001 to calculate industry-level employment growth for 2001.

A.1.6 PSID Estimation Timing: An Example

The data timing for the 2003 observation of a household interviewed in the first quarter of 2003 (2003 PSID wave) is as follows:

$$c_{2003}^{i,j} = \beta_0 + \beta_1^k y_{2002}^{i,j} + \beta_2^k BH_{2002q1}^j + \beta_3^k HPG_{2002q4,2001q4}^j + \beta_4^k BEG_{2002q4,2001q4}^j + \beta_5^k HY_{2001}^{i,j} + \beta_6^k FY_{2001}^{i,j} + \gamma^k \mathbf{X}_{2003}^{i,j} + \delta_{2003}^k + \eta^i + \varepsilon_{2003}^{i,j}. \quad (1)$$

That is, annualized consumption reported in the first quarter of 2003 by household i living in location j is estimated to be a function of household i 's income in 2002, bank health in the first quarter of 2002 in location j , local real house price growth and local employment growth between the fourth quarter of 2001 and the fourth quarter of 2002 in location j , and housing wealth-to-income and financial wealth-to-income measures from the previous PSID survey (2001).

A.1.7 PSID Sample Construction

As noted in the main text, we base our analysis primarily on the characteristics of the household head, and we focus on household heads aged 18 to 64. If the household head changes over the course of our sample, we create a new household. We drop observations in a given year for married or cohabitating

⁶ We construct NAICS-consistent employment data at the national level using the same approach that we use at the local level.

households when *both* the head and spouse/partner are students or retirees—this includes households where *both* adult members are younger than 18 years old or older than 64. If the head and spouse/partner are out of the labor force for other reasons, then the household is still included in our sample. For single (non-cohabitating) households, these requirements apply to the status of the household head. With regard to timing, these sample restrictions are based on a household member’s status as reported in the survey year.

We also omit household observations with extreme temporary changes in house values, mortgage amounts, or financial wealth-to-income ratios, as well as those with loan-to-value ratios greater than 2. We identify extreme temporary jumps or declines for one of the aforementioned variables as increasing by a factor of four or declining by three-quarters—changes that are reversed in the subsequent period conditional on the household not moving. Since households with small remaining mortgages may choose to pay them down, we exclude from these restrictions large declines in mortgage debt when the initial mortgage balance is less than \$25,000.

After applying these sample criteria, we impose additional restrictions. First, we must observe a household for at least two (consecutive or nonconsecutive) waves from 1985 through 2015, so that we can control for household-specific fixed effects. We further require that households report being either a homeowner or a renter in a given wave. In addition, homeowners must report a non-zero house value and must not switch to being a renter for a single period without indicating that they moved.⁷ We apply the same moving criterion to renters who report that they switch to owning, and we drop observations where households state they are a renter but report a positive home value. The number of households dropped due to these homeownership restrictions is quite limited.

Finally, we exclude observations in the top and bottom 1 percent of food consumption, household income, or the wealth distribution in each year. We do not adjust the wealth-to-income ratios that we use in equation (1) for outliers themselves, but instead drop outliers from the numerator and denominator before constructing the ratios. For the location-specific data (RENPL10, HPG, and BEG), we winsorize the data at the 1st and 99th percentiles of the distributions in each year rather than omitting the outliers, because omission would eliminate all households in a given location for that year. This approach helps avoid the undue influence of outliers on our results.

⁷ Homeowners who switch to being a renter for consecutive periods during our sample are assumed to have moved, even if the moving variable indicates otherwise.

A.2 Calculating Age-Earnings Profiles: Additional Details⁸

We take a number of steps to clean the PSID data prior to calculating individuals' age-earnings profiles. This data cleaning falls into two main categories: (1) adjusting anomalies in an individual's reported age profile and (2) identifying non-earners and business owners. In particular, we exclude observations from the estimation sample when the individual is retired, a student, not in the labor force, or is a business owner.

Adjusting an Individual's Age Profile

In some situations in the PSID, an individual's reported age diverges from his or her standard (increasing) age profile. These anomalies are likely due to either recording or reporting errors. An example of such a situation is the following: The PSID reports an individual's age as 23 in 1990, 24 in 1991, 25 in 1992, 26 in 1993, 45 in 1994, and 28 in 1995. Given the time series of data, the observation for 1994 is an outlier and most certainly incorrect. In situations like this one, we recode any outliers to fit with the natural (linear) age progression given the other available data. Specifically, we would change 45 to 27 in 1994 in this example.⁹

There are also cases in which an individual appears to jump from one (linear) age path to another higher (linear) path without a corresponding passage of time. For example, over seven consecutive years an individual's age profile may appear as 19, 20, 21, 22, followed by 35, 36, 37. Sometimes the age path will revert to one that is in line with the original path. When an individual has multiple age paths, as in this example, we choose the one that is consistent with the majority of that person's observed age records and recode the age data accordingly. In the case above, the path would become 19, 20, 21, 22, 23, 24, 25. If the observations for an individual are split evenly between two age paths, for example 42, 43, 44 followed (consecutively) by 55, 56, 57, then we choose the first path we observe and make the other age data consistent with that path (for example, 42, 43, 44, 45, 46, 47). In addition, due to the survey structure of the PSID, the age of an individual cannot realistically decrease between years. Therefore, when we see a non-transitory decline in the age profile of the household head or spouse/cohabitant for which we are constructing an age-earnings profile, we assume that this represents a new age path, treat the profile just like any other non-standard age path, and apply the cleaning procedures discussed above. If, instead, the individual's reported age declines in one year and then reverts to its original path, we adjust only the

⁸There are some nuances to the procedures described in this section, and additional information is available from the authors upon request.

⁹Note that an age path of 19, 20, 20, 22 could be a correct, error-free age profile if an individual were interviewed late in one year when he or she was 20 and then early in the next year when he or she was still 20. However, we smooth through these bumps to eliminate any unnecessary noise. So this path would become 19, 20, 21, 22.

temporary decline.

All of these adjustments leave us with realistic age paths for individuals in the PSID, and we use these revised age data where appropriate in calculating individuals' age-earnings profiles. That is, we enforce the appropriate age path, as needed, for each individual in our analysis.

Addressing Non-earners and Other Sample Restrictions

Including individuals in our age-earnings estimates who have no income or who have particularly volatile income from year to year (age to age) would likely result in less accurate age-earnings profiles for a given gender and education group. Therefore, we exclude certain individuals—mainly non-earners, but also business owners—when we estimate our group-based age-earnings profiles (equation [2] in the main text). In particular, we exclude any observation where the individual is a business owner, student, retiree, or is out of the labor force at a given age. All such individuals will have little or no labor earnings and including them would bias down our estimates of the group-based age-earnings coefficients.¹⁰ We identify these individuals in the PSID as follows:

- **Students:** Individuals who are younger than 18 years of age or who respond that they are a “student” when asked about their current employment status.
- **Retirees:** Individuals who are over 64 years old or who report their employment status as “retired” *and* whose labor income is 20 percent or less of what it was when they last reported their status as working. If individuals meet either of these criteria for at least five years (four years after 2000, when the survey becomes biennial) they are assumed to be retired for the remainder of the years we observe them in the data, regardless of their reported income. If individuals meet the retirement criteria for only one year, unless that year is the first or last year we observe them in the data, we assume that they are *not* retired and include them in our analysis.
- **Labor Force Nonparticipants (other than students/retirees):** We consider individuals to be out of the labor force if their labor income is zero or their reported labor income drops to 20 percent or less of its value in the previous period, and they do not report an unemployment spell. Individuals whose maximum income is \$5,000 or less over the entire period we observe them are also considered out of the labor force for all of their observations, including ones in which they report being unemployed, as are individuals who claim to be consistently unemployed and never report positive earnings.¹¹

¹⁰ Business owners' income often derives from capital and not labor, and any earnings that they have likely fluctuate a good deal from one year to the next, depending on the success of their business.

¹¹ This is likely a situation in which unemployment is reported or recorded incorrectly.

Individuals who appear out of the labor force for only one year, and that year is not at the beginning or end of the period over which we observe them, are coded as in the labor force.

- **Business Owners:** We identify business owners as individuals with nonzero business-asset income. In addition, if individuals report no business income for one period in a series of periods when they are business owners (and that period is not at the beginning or end of their participation in the PSID), we assume that they are business owners in all periods.¹²

We treat individuals who report that they are unemployed somewhat differently.¹³ In particular, we include them when we calculate the overall (group-based) age-earnings profiles (equation [2]), since experiencing a spell of unemployment does not mean an individual had zero earnings for a given year (age). Periods of unemployment are also a natural, but infrequent, part of a person's earnings profile over time. The incidence of unemployment and periods of lower-than-normal earnings likely varies by education group and age in ways that should be accounted for by our age-earnings estimates. In addition, as we note in the main text, we do not count periods of unemployment toward the minimum of five observations necessary to estimate an individual's person-specific age-earnings profile. However, we include periods with unemployment spells if at least five other observations are available where the individual is employed.

A.3 MSA-level Loan Growth Analysis

Data

We calculate real per capita non-mortgage loans in MSAs using the Federal Reserve Bank of New York Consumer Credit Panel (CCP). The CCP is a 5 percent random sample of all individuals in the United States with credit records from Equifax; the data begin with information from 1999. To be included in the Equifax data, individuals must have a social security number and have applied for credit at some point in their lives.

The CCP contains data on individuals' mortgage and non-mortgage credit. The non-mortgage credit data include auto finance, auto bank, bankcard, consumer finance, retail, and other loan balances. Because the CCP data on mortgage and home equity loan balances include inaccuracies, we supplement the main CCP data with information from the "Mortgage Tradeline" (MT) dataset, which is compiled from the Equifax data by staff at the New York Fed. The MT dataset is a loan-level database with information on as many as 14 of an individual's most recent first mortgages, five of his or her most recent home equity installment loans, and

¹² Prior to 1992, the income data are reported together for the household head and his or her spouse/partner, so we mark each as a business owner if this combined business income variable is nonzero.

¹³ We classify individuals as unemployed if they report one week (or more) of unemployment in a given period.

five home equity revolving loans. Since the MT database focuses exclusively on housing-related loans, it is restricted to current or former homeowners—individuals who at some point applied for a mortgage loan. The MT data improves upon the mortgage information in the main CCP dataset because sometimes small first mortgages are miscoded as home equity installment loans or large home equity loans are miscoded as first mortgage loans. The New York Fed properly classifies these loans in the MT data using information about loan origination dates, credit limits, or previous (maximum) balances. We present results using data only on nonmortgage loans as well as results where we add home equity loan data from the MT database.

We generate MSA-level loan data by assigning individuals to MSAs based on their reported zip code. When a zip code spans more than one MSA, we assign the individual to the MSA with the highest percentage of addresses within the zip code.¹⁴ Individuals with invalid zip codes are dropped from the calculations.¹⁵ The nonmortgage loan amount or nonmortgage plus home equity loan amount for a given MSA and quarter is the sum of all loans outstanding in that MSA and time period.¹⁶ Loans marked as “joint” are divided in half to avoid potential double counting.¹⁷ Per capita loans are calculated as total balances in an MSA divided by the number of people in the MSA in our CCP sample each quarter. We average these values across quarters to obtain annual loan values by MSA, which we convert to real values using the PCE deflator. We then calculate real, annual MSA-level per capita loan growth rates.

Estimation Equation

We analyze the relationship between bank health and MSA-level loan growth by estimating equations of the form:

$$g_t^m = \gamma_0 + \gamma_1 BH_{t-1}^m + \gamma_2 HPG_{t-1}^m + \gamma_3 BEG_{t-1}^m + \eta^m + \delta_t + \epsilon_t^m, \quad (2)$$

where g_t^m is real per capita loan growth (nonmortgage or nonmortgage plus housing equity) between year $t - 1$ and t in MSA m ; BH_{t-1}^m is our (baseline) measure of bank health (RENPL10) as of $t - 1$ in MSA

¹⁴ Typically, zip codes that span MSAs have addresses that are highly concentrated in one MSA compared with the other(s). However, there are some more ambiguous cases where addresses in a zip code are split about evenly across MSAs. In these cases, we randomly choose the MSA to which we assume the zip code belongs. The data on residential address ratios comes from the US Department of Housing and Urban Development’s Office of Policy Development and Research.

¹⁵ We assume that an individual’s zip code is correct, even if it appears inconsistent with the individual’s other location information (county and/or state).

¹⁶ CCP loans coded as “Authorized Use,” “Terminated,” or “Co-maker” are excluded from the MSA totals. We also drop loans in the MT data that are coded as anything but mortgage or home equity. While the MT data consist only of mortgage and home equity loans as designated by Equifax in the CCP data, the New York Fed staff reclassifies some of these loans as a different type.

¹⁷ A loan that lacks a joint/not joint classification is assumed to be held by the individual alone. In the MT data, a loan is assumed to be jointly held if its Equal Credit Opportunity Act (ECOA) code is “Joint Account” or “Shared.”

m ; BEG_{t-1}^m is the lagged one-year percent change in BEG in the MSA; and HPG_{t-1}^m is the lagged one-year percent change in (real) MSA-level house prices. The specification also includes year, δ_t , and location, η^m , fixed effects. To avoid undue influence on our estimates from outliers, we winsorize the top and bottom 1 percent of the bank health, loan growth, house price growth, and BEG distributions by year.

Table A1: Bank Health and Consumption: Initial Results
All Regressors

	(1) Food	(2) Food	(3) Food	(4) Food
Log Income: Owner	0.191*** (0.013)	0.191*** (0.013)	0.221*** (0.013)	0.223*** (0.013)
Log Income: Renter	0.242*** (0.013)	0.242*** (0.013)	0.296*** (0.015)	0.297*** (0.015)
RENPL10/100: Owner		-0.632 (0.424)	-1.534*** (0.486)	-1.550*** (0.486)
RENPL10/100: Renter		-0.866 (0.638)	-2.946*** (1.016)	-2.945*** (1.015)
Home Equity/(Inc.*100): Owner	1.176*** (0.387)	1.166*** (0.388)	1.003** (0.425)	0.832** (0.421)
Fin. Wealth/(Inc.*100): Owner				0.979*** (0.189)
Fin. Wealth/(Inc.*100): Renter				2.370*** (0.855)
Renter Dummy	-0.759*** (0.181)	-0.759*** (0.180)	-0.758*** (0.231)	-0.734*** (0.229)
House Price Growth: Owner	0.064 (0.062)	0.032 (0.060)	0.042 (0.061)	0.042 (0.062)
House Price Growth: Renter	0.059 (0.101)	0.018 (0.097)	-0.056 (0.120)	-0.057 (0.120)
Bartik Empl. Growth: Owner	0.013 (0.584)	-0.049 (0.586)	-0.051 (0.664)	0.001 (0.667)
Bartik Empl. Growth: Renter	-1.234 (1.069)	-1.252 (1.072)	-2.288 (1.605)	-2.187 (1.613)
Age: Owner	0.026*** (0.003)	0.026*** (0.003)	0.037*** (0.004)	0.037*** (0.004)
Age ² /1000: Owner	-0.309*** (0.031)	-0.308*** (0.030)	-0.413*** (0.040)	-0.420*** (0.040)
Family Size: Owner	0.190*** (0.015)	0.190*** (0.015)	0.171*** (0.018)	0.170*** (0.018)
Family Size Squared: Owner	-0.014*** (0.002)	-0.014*** (0.002)	-0.012*** (0.002)	-0.012*** (0.002)
Single Earner: Owner	-0.001 (0.022)	-0.001 (0.022)	-0.001 (0.027)	-0.004 (0.027)
Dual Earner: Owner	0.016 (0.022)	0.016 (0.022)	0.017 (0.025)	0.016 (0.025)
Age: Renter	0.038*** (0.005)	0.038*** (0.005)	0.042*** (0.006)	0.042*** (0.006)
Age ² /1000: Renter	-0.480*** (0.055)	-0.480*** (0.055)	-0.522*** (0.072)	-0.525*** (0.072)
Family Size: Renter	0.157*** (0.018)	0.157*** (0.018)	0.168*** (0.019)	0.168*** (0.019)
Family Size Squared: Renter	-0.014*** (0.002)	-0.014*** (0.002)	-0.019*** (0.002)	-0.019*** (0.002)
Single Earner: Renter	-0.028 (0.024)	-0.027 (0.024)	-0.079** (0.035)	-0.080** (0.035)

Dual Earner: Renter	0.024 (0.021)	0.024 (0.021)	0.008 (0.031)	0.008 (0.031)
Owner Obs.	53188	53188	25991	25991
Renter Obs.	39888	39888	19096	19096
Total Obs.	93076	93076	45087	45087
Adj. R-squared	0.593	0.593	0.573	0.574

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Dependent variable is log real food consumption. See Tables 2 and 3 for variable definitions. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A2: Bank Health and Consumption: Controlling for Constrained Households
Financial Wealth

	Negative Income Deviation		$\geq 5\%$ Neg. Dev.		Neg. Dev. or < 25		$\geq 5\%$ Neg. Dev. or < 25	
	(1) Unconst.	(2) Const.	(3) Unconst.	(4) Const.	(5) Unconst.	(6) Const.	(7) Unconst.	(8) Const.
Log Income: Owner	0.190*** (0.014)	0.191*** (0.015)	0.191*** (0.016)	0.192*** (0.016)	0.193*** (0.014)	0.193*** (0.015)	0.194*** (0.016)	0.195*** (0.016)
Log Income: Renter	0.233*** (0.020)	0.227*** (0.021)	0.235*** (0.022)	0.229*** (0.023)	0.237*** (0.020)	0.233*** (0.021)	0.241*** (0.022)	0.236*** (0.023)
RENPL10/100: Owner	-1.131** (0.524)	-2.420*** (0.621)	-1.409** (0.566)	-3.185*** (0.744)	-1.115** (0.523)	-2.428*** (0.622)	-1.394** (0.564)	-3.196*** (0.745)
RENPL10/100: Renter	-1.651 (1.029)	-3.584*** (1.051)	-1.680 (1.129)	-3.896*** (1.222)	-1.447 (1.044)	-3.669*** (1.035)	-1.447 (1.145)	-4.014*** (1.201)
Home Equity/(Inc.*100): Owner	1.170* (0.611)	0.639 (0.437)	1.057* (0.639)	0.430 (0.531)	1.190* (0.622)	0.629 (0.429)	1.084* (0.650)	0.423 (0.521)
Fin. Wealth/(Inc.*100): Owner	0.788*** (0.249)	1.057*** (0.253)	0.873*** (0.260)	1.213*** (0.311)	0.792*** (0.249)	1.063*** (0.250)	0.878*** (0.260)	1.219*** (0.307)
Fin. Wealth/(Inc.*100): Renter	-0.391 (1.269)	5.118*** (1.220)	-0.460 (1.319)	5.683*** (1.496)	-0.685 (1.348)	5.156*** (1.196)	-0.794 (1.394)	5.736*** (1.453)
Owner Obs.	21769	21769	19085	19085	21769	21769	19085	19085
Renter Obs.	12039	12039	10707	10707	12039	12039	10707	10707
Total Obs.	33808	33808	29792	29792	33808	33808	29792	29792
Adj. R-squared	0.553	0.553	0.547	0.547	0.553	0.553	0.547	0.547

Sources: Authors' calculations using PSID, Call Reports, FDIC Summary of Deposits, CoreLogic, and Bureau of Labor Statistics' Quarterly Census of Employment and Wages data.

Notes: Financial wealth sample. Dependent variable is log real food consumption. Definitions of constrained and unconstrained households are the same as in Table 4. See Table 2 for additional variable definitions. Specifications also include the following additional control variables (with separate estimated coefficients for homeowners and renters where applicable): local (Bartik) employment growth, local house-price growth, an indicator for whether the household is a renter, age, age squared, family size, family size squared, and indicators for single-earner and dual-earner households as defined in Table 2. Location (with separate estimates for homeowners and renters), year, and household fixed effects are also included. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A3: MSA Loan Growth Analysis: Summary Statistics

Variable	Mean	Median	Std. Dev	Min.	Max.	N
Excluding HE: Real Per Capita Nonmortgage Loan Growth (%)	-.27	-.29	6.1	-19	34	10,616
Including HE: Real Per Capita Nonmortgage Loan Growth (%)	1.7	.87	7	-17	36	10,616
RENPL10/100	.017	.0093	.017	.0015	.13	10,616
RENPL10/100 Multi-BHC (10%)	.016	.0083	.015	.00081	.093	9,606
House Price Growth (%)	2.8	3.2	6.9	-26	34	10,616
Bartik Empl. Growth (%)	.45	1.1	1.9	-7.5	3.2	10,616

Sources: Authors' calculations using the NY Fed Consumer Credit Panel/Equifax (CCP), Call Reports, FDIC Summary of Deposits, CoreLogic, and Quarterly Census of Employment and Wages data.

Notes: Annual data at the Metropolitan Statistical Area (MSA) level. Sample restricted to MSAs found in the PSID data. Sample that includes home equity is based on the subset of the CCP individuals with mortgage loan data (see main text for more details). RENPL10 is the baseline balance sheet (nonperforming real estate loan ratio) measure of bank health at the MSA level (see Section 2.2). Multi-locational banks are bank holding companies with less than 5 percent of their overall deposits held in that location, but where those deposits account for at least a 10 percent share of all local deposits.