The Effects of Changes in Local Bank Health on Household Consumption

Daniel Cooper and Joe Peek

Abstract:

This study investigates the relationship between credit availability and household consumption using a novel approach to separate credit demand and supply. We find that a deterioration in local bank health reduces household consumption, with the strongest effects occurring for households that are more likely to need credit—especially those experiencing a negative income shock and having limited liquid assets. The main contributions of the study are the use of an arguably exogenous measure of local bank health and multifaceted indicators of constrained households. Our findings contribute to the discussion of the linkages between the financial sector and real economic activity.

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.

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: February 2020; original version posted October 2018

1. Introduction

The 2007–2009 US financial crisis and the associated Great Recession have stimulated much research 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, with the literature primarily focusing on effects operating through household wealth (for a review, see Cooper and Dynan 2016). Indeed, oft-cited studies analyzing the impact of the financial crisis on consumption concentrate on the role of plunging house prices or the related topics of household financial distress (e.g., mortgage delinquency) and deleveraging (see, for example, Mian, Rao, and Sufi 2013; Dynan 2012; Albuquerque and Krustev 2018). However, the linkages between the financial sector and household consumption are broader and more complex. Both the availability of credit, mortgage as well as nonmortgage, and access to that credit are important determinants of household consumption, with many households relying on credit from banks with a local presence to help finance their expenditures. Thus, fluctuations in local-bank health, and thus local credit availability, which deteriorated during the financial crisis, can play an important role in determining the ability of households to maintain their desired level of consumption in the face of income and expenditure shocks.

This study explores the effect of local credit availability on household consumption beyond the typical drivers of household expenditures (a household's income, wealth, and demographic characteristics). Using an arguably exogenous proxy for credit availability based on the health of banks with operations in a locality (local-bank health), we find that, all else equal, households' (food) expenditures decline when local-bank health deteriorates. This effect is

particularly strong for households identified as "constrained" and potentially most in need of credit.[1]

While such results are not particularly surprising on their face, the contribution of this study arises from its novel method to address the major difficulty with determining the relationship between household consumption and household credit availability: isolating the credit supply from the credit demand effects. A further contribution arises from our multifaceted approach to identifying constrained households that includes an innovative method for capturing negative household income shocks and also considers the potential for households to insulate their consumption from income and expenditure shocks using their liquid assets (cash, stock and bond holdings). Faced with an adverse shock, households with limited liquid assets are likely particularly sensitive to the availability of external credit.

We reach these conclusions using household-level data from the Panel Study of Income Dynamics (PSID), Consolidated Reports of Condition and Income (Call Report) data on banks' balance sheets, and bank branch deposit data from the Federal Deposit Insurance Corporation (FDIC)'s Summary of Deposits (SOD) database. Since the condition of banks' balance sheets is affected by the local economy, we focus our analysis on the health of banks that are important to a given location but for which the location accounts for a small portion of the bank's overall deposits. Local economic conditions should therefore have little effect on the health of these "multi-locational" (ML) banks, making changes in their health arguably exogenous to local economic conditions.

Measuring local bank health based on the nonperforming loan (NPL) ratios of these ML banks, we relate credit availability to consumption after identifying constrained households that

---

[1] Throughout the study, we use the terms "liquidity constrained," "borrowing constrained," and "constrained" to reference households that are more likely to need or want access to credit in order to consume their desired amount.

are likely most in need of (bank) credit. In particular, we focus on households experiencing negative deviations in their current income relative to the level predicted by household-member-specific age-earnings profiles, as well as households' degree of income variability. Indeed, one might reasonably expect that even in the absence of large shocks, households with more variable incomes, on average, would rely more on credit to smooth their consumption, and thus be more sensitive to a deterioration in bank health. An additional dimension of credit needs arises from many younger households having incomes that are low relative to their expected future incomes. Such households tend to borrow in an effort to smooth their lifetime consumption, and thus also are likely to be disproportionately impacted by reduced bank credit availability. Income-constrained households with limited liquid assets (see, for example, Zeldes 1989) are also likely affected by credit availability. Indeed, liquid assets can help a household insulate its consumption from the effects of a negative income shock, or even an unexpected expenditure shock such as a large medical or car repair expense. However, if a household has few liquid assets, it is more likely that it will need to rely on external credit to smooth its expenditures over time and remain on, or near, its desired consumption path.[2] Unfortunately, a significant share of households have very limited liquid asset holdings, so that even relatively small income or expenditure shocks can impact their consumption, especially when bank credit dries up.[3]

Importantly, the link between bank financing and household consumption goes well beyond an effect on homeowners through the availability of mortgages and home equity lines of credit. Households in general rely on banks for personal loans, automobile loans, vacation loans,

---

[2] Using Canadian household-level data to examine the relationship between bank health and consumption, Damar, Gropp, and Mordel (2014) find that a reduction in credit supply reduces consumption for liquidity-constrained households, but not for households with adequate liquid assets to help finance their spending.
[3] 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." (Board of Governors of the Federal Reserve System, 2018)

and for other "big ticket" spending needs, such as furniture and large appliances.[4] Thus, a deterioration in bank health that reduces banks' ability and/or willingness to make credit available potentially affects a wide range of consumers, with income- or liquidity-constrained households likely impacted the most by a decrease in credit availability.

For our analysis, we estimate a standard consumption function augmented with a measure of local-bank health using data from the 1984 through 2015 waves of the PSID. We exploit both cross-sectional and time-series variation in our arguably exogenous measure of ML 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 households' real (food) consumption. We are careful to identify local banking options available to a household by focusing on bank branches operating in the household's location. Whereas the first-mortgage and credit card markets tend to be more national in scope, smaller local banks and local branches of larger banks likely play a more important role in providing home equity and consumer loans.[5]

Overall, we find that a deterioration in local-bank health reduces households' consumption, even after controlling for fundamentals that typically predict household expenditures. Importantly, the results show that the consumption of constrained households is affected more strongly than the consumption of unconstrained households. Indeed, the reduction in real food consumption based on a deterioration in bank health equivalent to moving from the 25[th] to 75[th]

---

[4] 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 often comes in part from banks. Thus, bank financing likely continues to play an important, albeit more indirect, role in consumer-credit availability.
[5] In the context of small-business lending, Nguyen (2019), among others, shows that even after the improvements in information technology, distance from bank branches played 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 consumer 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.

percentile of the bank health distribution is equal to the effect of a 1.5 percent decline in income relative to the mean for constrained households compared with a 0.6 percent decrease for unconstrained households. In addition, the households whose consumption is affected the most by a deterioration in bank health are those that experience a negative income shortfall or are young *and* have limited liquid assets. These households lack the internal resources to maintain their desired level of consumption without access to external funds. We further demonstrate that our effects are strongest in locations with the worst bank health—consistent with the weakest banks coming under binding regulatory constraints that limit their ability to supply credit. In addition, our findings are robust to alternative approaches to measuring bank health and to employing broader measures of household consumption.

This study 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, Peek and Rosengren 2000; Chava and Purnanadam 2011; Bassett et al. 2014). More specifically, our results provide a 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 studies include Abdallah and Lastrapes (2012), which examines the response of retail 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 collateralized by 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 studies, our research is generally broader; we are interested in the impact of changes in the availability of local-bank credit on household spending, highlighting the effects on constrained households, both due to negative income shocks and due to liquidity constraints.

The remainder of the study 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 results, and Section 6 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, as it is the only spending measure available in nearly all survey waves. Financial asset and non-housing debt data are available in the PSID surveys for 1984, 1989, and 1994, and from 1999 onward. In addition, the PSID does not contain information on the identity of the specific bank(s) from which a household obtains credit; instead

we link households with the set of banks that operate in their local market. Section A.1.1 in the Appendix has more details about the PSID data and linking households with local banks.

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 household over time and control for household-specific spending habits.[6] In addition, the PSID contains comprehensive data on household wealth and income, and, importantly, 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 the CEX—allows us to link households to local credit and economic conditions.

2.2 Bank Health Data

Because we want to capture local-bank health for locations with multiple banks, we compute a location-specific measure based on the health of individual banks operating in that location. In particular, we combine balance sheet information from banks' Call Reports with individual bank-branch deposit data from the FDIC's SOD. The SOD data are important for calculating local bank health since, 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. We use the Call Report and SOD data to generate measures of bank health at the local (MSA) and state rural (non-MSA) levels.[7] (Section A.1.2 in the Appendix provides further details about how we generate our sample of local banks.)

---

[6] The CEX has a panel dimension, but it is very short (four quarters) compared with the PSID.
[7] 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. We convert the annual SOD data to a quarterly frequency to match the quarterly Call Report data by using the annual SOD deposit value for each quarter in that calendar year.

Our baseline measure of bank health is based on banks' NPL ratio, defined as loans past due 90 days or longer and still accruing, plus nonaccrual loans, divided by total bank assets. We focus on the health of the bank branch at the median of the NPL distribution in each location in each period. We identify this bank by ordering all bank branches in a given location by their NPL ratio, from poorest health (highest NPL value) to best health (lowest NPL value). We then cumulate local bank deposits, starting with the bank branch that has the highest NPL ratio (worst health). The NPL ratio for the bank branch at the median of the local deposit distribution (NPL50) is our measure of bank health for that location and period.[8]

We base our bank health measure on total NPLs rather than on real estate NPLs even though 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. However, total NPLs provides a more comprehensive measure of the state of banks' balance sheets. Our results are quite similar, however, if we use an analogously constructed bank health measure based on the median of the real estate NPL ratio distribution in each location and year. In addition, we prefer basing our bank health measure on a NPL ratio instead of another bank balance sheet measure, such as a capital-to-assets ratio, because capital-to-asset ratios exhibit systematic differences across bank asset-size classes.[9]

Multi-Locational Bank Health

---

[8] For example, assume that a given location in a certain year has total bank-branch deposits of $100 million. First, we order the bank branches in that location and year by the NPL ratio of their associated parent bank. We then start with the bank branch with the worst NPL ratio and cumulate deposits until we reach a total of $50 million. The bank branch accounting for that cumulated deposit level is identified as the bank branch at the median of the NPL distribution, and we record the NPL ratio of that branch's parent bank as our NPL50 value for that location.

[9] A bank's willingness to lend is also likely more closely tied to the difference between its desired and actual capital ratios.

Although we control for local economic activity in our analysis (see Section 2.3), the potential endogeneity of our local-bank health measure remains a concern, given that the health of the local economy likely affects the health of banks operating in that location. 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 in a particular location and year will have little measurable effect on the health of a ML banking organization, whose operations and deposits are spread over many locations, if the location accounts for very little of its activity. We therefore focus our analysis on the health of ML banks, since such a measure is arguably exogenous with respect to local economic conditions.

We designate a bank branch in a given location as ML if all the branches of its bank holding company (BHC) in that location account for 5 percent or less of the BHC's total deposits across all locations. This requirement ensures that the location does not represent a meaningful part of the BHC's overall operations. In addition, for the location to be included in our analysis, the total deposits at all ML bank branches operating in that location must account for at least 5 percent of the overall deposits in that location to ensure that the ML bank branches are relevant for local credit supply. Section A.1.3 in the Appendix provides additional details about how we define ML banks.

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 rural level for households residing in non-metro areas. In particular, we control for local (real) house price growth (HPG) and local employment growth. Since HPG is location specific, and not tied to homeownership status, it can serve as a 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 core-based statistical area (CBSA, commonly known as MSA) and state levels.

We also condition our estimates on a Bartik-style (Bartik 1991) measure of local employment growth to absorb fluctuations in local business cycle conditions. This measure is constructed by combining data on national employment growth by industry with Quarterly Census of Employment and Wages (QCEW) data on employment by industry 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 local employment growth. As with HPG, we calculate BEG at both the MSA and state level, assigning state-level BEG to rural PSID households. (See Sections A.1.4 and A.1.5 in the Appendix for further details about how we construct HPG and BEG, respectively.)

2.4 External Validity: ML 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 the Federal Reserve Bank of New York Consumer Credit Panel (CCP) provided by Equifax. Including only the set of MSAs for which we have constructed our bank health measure, we regress annual real loan growth on our (one-year) lagged baseline measure of bank health (ML NPL50), as well as on controls for lagged local economic conditions (HPG and BEG). We also include year and location fixed effects—see Section A.2 in the Appendix for further details about this analysis and the CCP data.

The regression shows a strong, statistically significant, negative correlation between ML NPL50 and real per capita nonmortgage loan growth (see Table 1, column 1), all else being equal. Since high levels of ML NPL50 indicate poor bank health, the negative estimated coefficient is consistent with real loan growth declining as bank health deteriorates. When we include home equity loans (but not primary mortgages) in our measure of loan growth (column 2), which also tend to be originated by local banks, the estimated negative effect quadruples and is still highly significant. Overall there appears to be a meaningful (negative) correlation between our measure of local bank health and 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, household wealth relative to income, and a series of 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:

$$c_t^{i,j} = \beta_0 + \beta_1 y_{t-1}^{i,j} + \beta_2 BH_{\hat{t}-1}^j + \beta_3 BH_{t-1}^j \cdot V^i + \beta_4 HPG_{\hat{t},\hat{t}-1}^j \quad (1)$$

$$+ \beta_5 BEG_{\hat{t},\hat{t}-1}^j + \beta_6 WY_{t-1}^{i,j} + \gamma X_t^{i,j} + \delta_t + \eta^i + \epsilon_t^{i,j},$$

where $c_t^{i,j}$ is the logarithm of annual real food 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, and V^i is a measure of household i's income volatility. $HPG_{\hat{t},\hat{t}-1}^j$ and $BEG_{\hat{t},\hat{t}-1}^j$ are appropriately timed house price

11

growth and (Bartik) employment growth in location j. $WY_{t-1}^{i,j}$ is the lagged ratio of household

wealth relative to income (WY). $\boldsymbol{X}_t^{i,j}$ is a vector of household demographic controls, including

age and age squared for the household head, family size, family size squared, homeownership

status, and whether the household includes multiple earners.[10] δ_t^k is a set of annual (1,0) dummy

variables (year fixed effects). η^i is a household-specific fixed effect. We interact income

volatility with bank health $(BH_{t-1}^j \cdot V^i)$ to capture that the expenditures of households with more

volatile income may be more sensitive to changes in credit availability.[11]

Our empirical approach is based on the reduced-form, hybrid log-level consumption

estimation equation proposed in Muellbauer (2008) and used in Cooper (2013). Controlling for

household wealth in addition to household income is standard practice when estimating reduced-

form consumption equations. Our measure of household wealth from the PSID includes the

current value of financial assets less any (nonhousing) debt for all households and adds housing

equity values for homeowners (house value less any outstanding mortgage debt).[12] In addition,

using a levels-based setup allows us to include PSID households in our analysis 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.

Our coefficients of interest are β_2 and β_3, with the former coefficient capturing the direct

effect of bank health on consumption and the latter coefficient capturing the incremental effect

---

[10] We include a (1, 0) dummy variable for whether the household includes a spouse/cohabitant but at most one earner, and a (1, 0) dummy variable for whether the household includes a spouse/cohabitant and there are two earners. The third category, single, is the excluded category.
[11] We define income volatility as the standard deviation of the yearly deviations in households' actual income from the level predicted by their (individual-specific) age-earnings profiles, which we describe in detail in Section 4 and Section A.2. Income volatility does not enter equation 1 directly, as it is absorbed by the household fixed effect.
[12] See Section A.1.1 in the Appendix for further details regarding the asset and debt data available in the PSID. Since household wealth data are not available in every PSID wave, we also provide estimates of equation (1) that exclude $WY_{t-1}^{i,j}$ and are able to use a much larger sample of households.

based on differences in income volatility across households. Estimates of β_2 that are different from zero imply that bank health has a direct, independent effect on consumption, even after controlling for local economic conditions and other demographic and financial factors that typically are important predictors of household spending. Similarly, estimates of β_3 that are different from zero imply that, all else equal, income volatility has a differential effect on the relationship between bank health and consumption.

Since our primary (ML NPL50) and alternative measures of bank health are constructed such that higher values represent banks that are more financially troubled, finding $\beta_2 < 0$ suggests that consumption falls when bank health deteriorates and rises when bank health improves. Similarly, finding $\beta_3 < 0$ would imply that for a given level of bank health, greater income volatility, and hence a potentially greater need for credit, leads to lower consumption (presumably because households cannot get the credit they need to smooth through their income fluctuations) and a bigger overall bank health effect. We further anticipate that the estimated bank health effects will be stronger (in absolute value) for constrained households.

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 a household's information set at approximately the time it makes its consumption decision, given the timing of the available data in the PSID. Variables dated t in equation (1) are measured for the quarter in which each household's PSID interview occurs, while variables dated $t\text{-}1$ cover the year prior to the survey year. For example, income data, which are recorded for the calendar year prior to the survey year in the PSID, are dated $t\text{-}1$. Household wealth relative to income is also measured as of $t\text{-}1$ to capture a household's wealth 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, the household wealth data are lagged two calendar years when constructing wealth-to-income ratios once the PSID switches to biennial data in 1999.[13] Demographic variables, such as age, family size, and homeownership, are measured at time *t*, while single versus dual-earner households are determined based on data from *t-1*. Our household constraint indicators are also based *t-1* data.

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) real 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 local-banks' health at the beginning of the period in which households make their annualized spending decisions. See Section A.1.6 in the Appendix for a concrete example of how we address these data timing issues 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. Including household wealth limits our sample, because, given our timing conventions, lagged wealth data are available to estimate equation (1) prior to 1999 only for 1984, 1989, and 1994 (the 1985, 1990, and 1995 regression sample years). 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, with lagged wealth and income data beginning in 1984, and includes households in which either the head or

---

[13] With the switch to biennial data, household wealth is scaled by average household income from the years on either side of the year wealth is measured to approximate household income in the wealth year. For example, wealth in 2009, which we pair with 2011 consumption, is scaled by average family income from 2008 and 2010.

spouse/cohabitant (if any) is at least 18 years old and no older than 64.[14] Households must also

have non-missing data for all demographic variables in equation (1), including non-missing

wealth data where applicable. Our final estimates include 29,431 observations for the wealth

sample (54,109 observations for the full sample that excludes WY from equation 1). The wealth

sample includes 6,347 unique households (8,194 unique households for the full sample).

4. Identifying Constrained Households

Overview

Whereas the Survey of Consumer Finances (SCF) contains direct questions about

households' access to credit, most datasets, including the PSID, do not, and information on

whether households are liquidity constrained must be inferred. Indeed, some studies, 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. Other approaches identify

potentially constrained households within the PSID based on the age of the head of the

household (Lehnert 2004) or liquid asset holdings (Zeldes 1989). However, without direct credit

constraint questions, there is not necessarily a one-size-fits-all approach to determining whether

a household is constrained, and using multiple approaches could be beneficial.

While we employ age and liquid asset (relative to income) information to identify

potentially constrained households, we also 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 the percent deviation in households' current income relative to

the level of income predicted by age-earnings profiles with household-member-specific

---

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

intercepts. This approach, which we describe in detail in Section A.3 of the Appendix, allows us to identify years in which households likely face negative income shocks.[15]

Importantly, we view these negative household income deviations as transitory. That is, our approach is designed to identify temporary income shortfalls (windfalls) rather than more permanent income declines (increases) that would lead households to permanently reduce (increase) their expenditures. Indeed, the data suggest that our estimated income deviations are relatively transitory. Only about 40 percent of households that experience a negative income deviation in one period have another negative income deviation in the subsequent period, and only 6 percent of households experience negative deviations for three consecutive periods. These probabilities are slightly less than what one would expect if the subsequent negative shocks were random conditional on experiencing a negative shock in the first period, consistent with our view that the shocks are transitory. (The pattern of the share of households experiencing consecutive positive shocks is qualitatively similar, although the shares themselves are slightly larger.)

Determining Likely Constrained Household Observations

We use our measure of income deviations along with other data to identify potentially constrained households in the PSID. First, any household with a negative income deviation in a given period is deemed to be constrained in that period. We further refine this indicator by combining it with data on the age of the earner(s) in the household. Even in the absence of a negative income deviation, households that are relatively young are more likely to demand credit in an effort to raise their current consumption toward their lifetime average, consistent with the

---

[15] 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.

implications of the Life Cycle Hypothesis. An additional constraint indicator captures households' ability to insulate their consumption from negative shocks to their income through relying on their liquid asset holdings. In particular, we designate households with below median liquid assets relative to income (LA) as being constrained.

Our most comprehensive measure of potentially constrained households combines information on their age, income deviations, and LA. Arguably, households with lower relative liquid asset holdings—especially those households that are young—that experience negative income shocks have less ability to use liquid assets to finance their desired consumption and are thus more likely to demand credit. If such credit is less available, or even unavailable, these households are more likely to have to curtail their current consumption expenditures. Therefore, these households' consumption should be most susceptible to a deterioration in local-bank health that results in a contraction in credit availability.

5. Results

Overall, we find a negative relationship between bank health and household spending. As banks' health deteriorates, they tend to make less credit available, resulting in reduced consumption expenditures for those households that rely on credit to help finance their expenditures. This bank health effect is the strongest for households that we identify as constrained—especially those households with income shortfalls and little, if any, LA. We further demonstrate that credit availability is a driving force behind our findings, as consumption suffers the most in locations with the worst bank health.

Table A.1 in the Appendix provides summary statistics for the baseline sample we use to conduct our analysis and reach these conclusions. In general, the data look reasonable. For the full sample where we do not control for WY (left columns), mean household income (2000

dollars) is about $49,000, mean real food consumption is around $5,500, and the share of homeowners is 64 percent. The statistics for the full sample and the wealth sample (right columns) are quite similar, suggesting that our two samples are not vastly different based on observables. The somewhat higher income and expenditure values for the wealth sample can be attributed to it being drawn disproportionately from more recent household observations. Importantly, the bank health data are quite similar for the two samples, suggesting that the characteristics of banks' health are not sensitive to which sample we consider. In addition, Table A.2 in the Appendix shows the (unconditional) correlation matrix for our main regressors of interest. Importantly, the correlation between bank health and household income or wealth is quite low.

5.1 Baseline Estimates

Our initial estimates of equation (1) use a measure of bank health based on the median nonperforming loan ratio for all banks with branches in a location (NPL50). These results show the expected negative relationship between bank health and real food consumption (Table 2, column 2)—an effect that is statistically significant at the 5 percent level. Interacting our measure of household income volatility with bank health also produces an (expected) negative coefficient that is significant at the 5 percent level. Thus, for households with more volatile income, the negative bank health effect on consumption is even stronger. This finding is consistent with these households having a greater need for (bank) credit to smooth their consumption over time, all else equal, and thus their expenditures are more sensitive to banks' health. Note that the memo line in this and other result tables shows the total effect of bank health on consumption evaluated at the mean of household income volatility to more easily compare the magnitudes of the overall bank health effects across specifications.

18

Including bank health in the analysis does not noticeably alter the estimated relationships between consumption and income or consumption and HPG. In contrast, the point estimate of the effect of local employment growth on consumption switches from a small positive to a small negative value, although such a shift is not particularly surprising given the imprecision of the estimates (compare Table 2, columns 1 and 2). For presentation purposes, we show estimates for only the most relevant regressors in Table 2, while Table A.3 in the Appendix shows results for the full set of regressors. The estimated effects for our other demographic control variables are consistent with what we would expect based on life-cycle consumption theory.

Because the overall measure of local-bank health used in column 2 is likely tied to local economic conditions, and thus the estimated bank health relationship is not necessarily exogenous despite our location-specific controls, in column 3 we switch to the equivalent measure of bank health for the subset of ML banks in a given location (ML NPL50)—a measure that is arguably more exogenous as we discussed in Section 2.2. These results are quite similar to those with NPL50 shown in column 2. While the direct and total effects on consumption are more precisely estimated using ML bank health, the total effect of bank health on consumption— evaluated at the mean of income volatility—is slightly smaller, consistent with the removal of endogeneity from the bank health measure. Given that ML bank health is arguably more exogenous, we use ML bank health measures for the remainder of our analysis.

When we focus on the (substantially smaller) sample for which we have household wealth data (columns 4 to 6), the direct bank health effects are somewhat larger (in absolute value), as are the total bank health effects that take into account income volatility.[16] However, the estimates show that much of the increase in the magnitude of the bank health effect comes from the change

---

[16] The estimated income volatility interaction effects are slightly smaller in this sample, but mean volatility is higher.

in the regression sample (column 4) rather than controlling directly for household WY (column 5). The results in column 6 further show that controlling separately for households' assets and liabilities, each measured relative to income, rather than for their overall WY has virtually no impact on the estimated bank health effects. Indeed, we cannot reject the null hypothesis that the estimated coefficients on the two ratios are equal and opposite signed (p-value 0.45), and thus we include households' overall WY going forward for ease of presentation.

Overall, our initial results suggest that bank health has an independent impact on household consumption beyond the effects of household wealth and income. For the remainder of our analysis, we focus on the specification in column 5 of Table 2.

5.2 Bank Health and Constrained Households

Initial Definitions of Constrained

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 the consumption of constrained households than for unconstrained households—a result that holds across different approaches for identifying constrained households.

Using a negative income deviation (shortfall) to designate potentially constrained households (Table 3, columns 1 and 2),[17] we find that a deterioration in bank health has a negative (direct) impact on consumption for both constrained and unconstrained households. However, the effect for constrained households is more than double the effect for unconstrained households (in absolute value) and is more precisely estimated. (The p-value for the test that the two estimated effects are the same is 0.048). The *total* effect of bank health on the consumption

---

[17] In Table 3 and elsewhere, we present estimates for constrained and unconstrained households in separate columns for ease of comparison. However, the effects for these households are estimated in the same regression so that the estimated coefficients for the location, time and household fixed effects, as well as the demographic controls, are unchanged as a household switches between being constrained or unconstrained from year to year.

of constrained households—taking into account the income volatility effect—is also more than twice as large as that for unconstrained households, and the difference is statistically significant (p-value 0.022). Indeed, households subjected to negative income shocks appear more dependent on local banks' ability and willingness to lend—that is, the availability of bank credit—in order to finance their expenditures and smooth their consumption over time. Moreover, the significance at the 10 percent level found for the (negative) direct and total bank health effects for unconstrained households may reflect a sensitivity to credit availability for these households associated in large part with expenditure shocks, such as unexpected large medical expenses. Faced with such a shock, even a relatively unconstrained household may have to reduce its other spending, including on food, if credit is not readily available.

The difference in the effect of bank health on consumption for constrained versus unconstrained households is more noticeable when we augment our negative income deviation definition for constrained households with information about the age of the household. The specifications in columns 3 and 4 of Table 3 incorporate young households in our definition of constrained, regardless of whether these households experience an income shortfall. In particular, we label a single-earner household as constrained if the worker is younger than 30 years old; a dual-earner household is constrained if both earners are younger than 30. 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. With this expanded definition, both the direct and total effects of bank health on constrained households' consumption are more than three times as large and are much more precisely estimated than the corresponding effects of bank health on unconstrained households' consumption. (Both

21

differences are statistically significant with p-values less than 0.003.) This finding is consistent with bank health mattering more for the consumption of households that are most likely to need credit to help finance their desired expenditures.

The results are similar if instead of grouping household observations based on their income deviations, we group them based on whether they have above-median LA (unconstrained) or below-median LA (constrained)—see Table 3, columns 5 and 6. Households in the below-median category have liquid assets that are 3.7 percent or less of their income, with about 50 percent of the group having no liquid assets. As a result, these households have limited ability to internally finance negative income shocks or expenditure shocks, and thus are likely to need external credit to achieve their desired consumption path. This alternative approach for identifying constrained households, which does not rely on household income shocks, yields a similar pattern. Constrained households' consumption is more sensitive to a deterioration in bank health than unconstrained households' consumption, with the total effect for constrained households being roughly 70 percent larger. If we further expand our LA-based definition of constrained to include young households (columns 7 and 8), the difference between the estimated effects is even larger, with the consumption of unconstrained households becoming less sensitive to bank health. Overall, bank health appears to matter most for the consumption of constrained households regardless of the specific definition used to determine such households.

Interpretation:

The point estimates for the bank health effects for constrained (and unconstrained) households in Table 3 (columns 3 and 4) imply that, all else being equal, going from the 25th to 75th percentile of the ML NPL50 distribution (1.4 percentage points) leads to roughly 3.3 percent lower annual real food consumption for constrained households relative to their average

expenditures and about 1.1 percent lower annual real food consumption for unconstrained households. These effects amount to about $184 less real food consumption for constrained households, compared with $68 less real food consumption for unconstrained households.

In absolute terms, bank health has a relatively small dollar impact on constrained households' real food consumption. Still, the effects are not trivial, as the estimated dollar decreases are equivalent to the impact on food consumption of a roughly 1.5 percent decline in income, relative to the mean, for constrained households, and a 0.6 percent decline for unconstrained households. These estimated effects may also be on the low side relative to the true impact of bank health on household consumption. In particular, most food expenditures tend to be relatively less discretionary than other spending categories (although the additional cost of eating out is discretionary), so households may be more likely to adjust their spending elsewhere in response to a decrease in credit availability. Our bank health estimates also capture effects that are averaged across households that fit our definitions of being constrained (or unconstrained), and the actual effects of more limited credit availability for individual households within these groups could be larger based on the extent of their income deviations and/or relative liquid asset holdings.

Combining (Negative) Income Deviation and Liquid Asset Data

When a household is subjected to a negative income shock, having access to liquid assets or to external credit becomes important for insulating consumption from the shock. We would therefore expect households that experience a negative income deviation *and* have low LA to be particularly in need of credit to finance their expenditures, and thus be most sensitive to bank health and credit availability.

Table 4 reports estimates where we combine our definition of constrained households based on negative income deviations and age with information on their liquid asset holdings. We find that *constrained* households with low (below median) LA have the largest (most negative) direct and total bank health effects (column 2). The consumption of constrained households with high (above median) LA also declines when bank health deteriorates, but not by as much, insofar as households in this group that experience negative income shocks have relatively larger liquid asset buffers that they can use to at least partially finance their consumption.

The consumption of unconstrained households with low LA also declines when bank health deteriorates, but to a lesser extent. Even though these households are not subject to a negative income shock, as we discussed earlier, these households could experience unexpected (nonfood) expenditure shocks that pressure their ability to maintain their level of food consumption in the absence of sufficient credit availability. Finally, the consumption of unconstrained households with high LA has by far the smallest (in absolute value) point estimate for the sensitivity to bank health—a result that is not surprising, given that these households should be the least likely of the four groups to need bank credit.

In addition, it is possible that the differences in the bank health effects between constrained and unconstrained households with different (relative) levels of liquid assets would be more apparent using alternative approaches for defining constrained households. However, we are limited by the data that are available in the PSID. Households with higher levels of liquid assets may also still be sensitive to bank credit availability because they choose to borrow rather than spend down their liquid assets in order to maintain a minimum level of liquidity for their possible future needs when they might be unable to borrow. Indeed, consistent with this type of hedging behavior, Gorbachev and Luengo-Prado (2019) show that households at times choose not to

repay high cost credit card debt in order to maintain or build a sufficient liquidity buffer against future shocks.

The relationship between credit needs and consumption is further highlighted in Table 4 by the finding that income volatility matters the most in the relationship between consumption and bank health for constrained households with low LA. This result is consistent with the idea that households with more volatile income likely have a greater need to finance their expenditures through other means, and those households with limited liquid assets are most likely to require credit to do so, and thus are even more sensitive to bank health. Overall, the results in Table 4 show that bank health impacts household consumption beyond the standard effects of income and wealth, especially for households most likely to require credit to finance their expenditures.

5.3 Mechanism

Even though our results fit with our priors, it is worthwhile to delve deeper into the possible underlying mechanism(s) driving our findings. The primary mechanism we have in mind is based on credit demand and credit supply. Households demand credit when they wish to finance expenditures beyond what their current income and liquid assets allow. This situation can occur when a household is subjected to a negative income shock, has low current income relative to expected future income, or suffers an unexpected expenditure shock, the latter of which we cannot explicitly account for with the available data. With regard to credit supply, the health of bank balance sheets and other (economic) factors drive banks' ability and willingness to provide credit to consumers and businesses. When banks make less credit available to consumers, then those households that wish to borrow to finance their expenditures may not be able to do so, and thus their consumption will fall relative to their desired spending path, unlike (or at least to a greater extent than) in times when credit is more readily available.

While distinguishing between credit demand and credit supply is a difficult task, we have tried to employ a proxy for bank credit availability that is arguably exogenous to the local economic conditions that may impact households' desired expenditures, and hence their demand for credit. In addition, just because banks' balance sheets become somewhat less healthy due to (in our case) an increase in nonperforming loans, it does not mean that lending will shrink substantially. For example, if a healthy bank becomes more (less) financially sound, likely easing (tightening) a nonbinding constraint, it is unlikely that the bank will sharply increase (decrease) loan supply. In contrast, if a weak bank's health deteriorates further (tightening a potentially already binding or near-binding constraint), it is more likely that the bank will substantially curtail its lending and not fully meet consumer loan demand. Therefore, if the mechanism we have in mind for explaining our results is valid, we would expect to see the greatest impact of bank health on consumption when local banks are in particularly bad shape.

Distinguishing among locations based on the distribution of bank health suggests that this is indeed the case. In particular, the relationship between bank health and consumption is strongest (and statistically significant) for households in locations where bank health is the worst. Table 5 presents results where we group locations (using (1, 0) dummy variables) based on quintiles of the overall bank health distribution. For example, "ML NPL50 Quint.5" is an indicator variable that takes a value of one for households living in locations where ML NPL50 is in the highest quintile (bank health is the worst) in a given year and is zero otherwise. The estimated coefficients are all expressed relative to the group of locations with the lowest (best) ML NPL50 readings (quintile 1, the excluded category). The estimates imply that the consumption of households living in a location that is in the worst quintile of the bank health distribution have about 5 percent lower food consumption, on average, than households living in

locations where banks are the healthiest. In comparison, consumption is only roughly 2 percent lower in locations that fall in the 2nd through 4th quintiles of the bank health distribution, and only the effect for quintile 5 (worst BH quintile) is statistically significant. These findings are consistent with banks with the very worst health curtailing their lending the most, and when this happens, households that need credit to finance their consumption that are living in locations with these unhealthy banks are forced to cut back on their spending relative to households living in locations with healthier banks.

In addition, Table 5 shows results where we do and do not interact income volatility with the bank health quintiles. The direct effects of bank health are nearly identical across these two specifications—consumption is lowest in locations where bank health is the worst. (We do not present total bank health effects in the table because the interaction terms are all small and statistically insignificant.) Overall, the estimates in Table 5 are consistent with our bank health measure serving as a proxy for credit supply, and with the effect of bank health on consumption being nar, as the cut back in credit availability is much more severe for those banks in the weakest tail of the bank health distribution.

This credit supply story is further reinforced by the results in Table 6, which show that consumption is most sensitive to bank health for constrained households living in locations with the worst bank health—especially constrained households with low LA. Since these households are the ones most likely to want to borrow, the fact that their consumption is affected the most when living in locations with the weakest banks, and thus much more limited credit availability, is consistent with a credit supply channel. There is even a sizeable effect for constrained households with low LA living in locations in the 4th quintile of the bank health distribution—a

result that likely reflects these households' need for credit in locations where it is limited, even if not the most limited.

We gain additional insight into the credit demand effects driving our results when we take a slightly alternative approach to defining constrained households based on their income deviations. Table 7 presents estimates from the same specifications as in Table 4, but defines income-constrained households as having a 5 percent or greater negative income deviation or being younger than 30 years old. Unconstrained households are defined the same as before (zero or positive income deviation and not younger than 30), so households with small negative income deviations (between zero and –5 percent) are excluded from this analysis. This approach yields little difference between the estimated direct bank health effects for constrained households with high LA compared with low LA. In comparison, when we defined constrained households based on any sized negative income deviation (Table 4), we found a larger difference—constrained households with low LA were much more sensitive to bank health than households with high LA. More specifically, the estimated direct bank health effect is roughly the same for constrained households with low LA in Table 7 as in Table 4, whereas the estimated effect for constrained households with high LA is larger (in absolute value) in Table 7.

It is perhaps not surprising to find a greater sensitivity to bank health for constrained households with high LA under the alternative constraint definition in Table 7, since the negative income deviations we consider are on average larger. Indeed, many of the income-constrained households with small negative income deviations but high LA, who are omitted from the analysis in Table 7, may have had sufficient liquid assets to insulate their consumption from these small shocks. Thus, their spending was not that sensitive to a deterioration in bank health. However, when we consider only the larger negative income deviations in Table 7, constrained

households with high LA may no longer have sufficient liquid assets to cover their income shortfalls. In contrast, the relationship between consumption and bank health is much more similar across specifications for constrained households with low LA likely because these households need credit to smooth their expenditures even for relatively small negative income deviations. Despite the different results for constrained households given the different income deviation definitions in Tables 4 and 7, the bank health effects for constrained households are still larger (in absolute value) than the effects for unconstrained households.

Taken together, the results in Tables 4 and 7 highlight the importance of households having sufficient liquid assets when experiencing a negative income shock. Indeed, when households have quite limited LA, negative income shocks can matter even when the shock is small, pushing these households to turn to bank credit in an effort to maintain their desired consumption path. Such households living in locations with worse bank health, and thus more limited credit availability, are therefore forced to cut back on their consumption by more than is the case when bank credit is more readily available. Overall, we believe the differences between the results in Tables 4 and 7 are indicative of the interplay between household income shortfalls, relative liquid asset holdings, and the availability of bank credit that is driving our results.

5.4 Robustness

Our results are robust to alternative ways of measuring bank health as well as consumption, as we discuss in detail in the Appendix (Section A.4). Alternative measures of bank health include ones based on the lower end of the ML NPL distribution as well as a measure of banks' real estate nonperforming loan ratio. We also consider a measure of bank health based on confidential bank supervisory (CAMELS rating) data and explore alternatives for how we determine ML banks. With regard to consumption, we consider alternative measures

based on augmenting real food consumption with additional, likely discretionary, spending data that are available in the PSID starting in 2005. We also utilize a broader (imputed) nondurable measure of consumption based on the approach in Blundell, Pistaferri, and Preston (2006).

In all instances the results, which we report in the Appendix (Tables A.4 to A.9), are qualitatively similar to our main findings in Table 4. It is particularly reassuring that we obtain results very similar to our main findings using the CAMELS data, which is a completely different approach to measuring bank health. These results suggest that our findings are not being driven by our decision to focus on banks' balance sheet data to capture their relative health. Finally, we show that our results are not simply picking up different, time-varying, macroeconomic factors across locations, as our estimates are robust to including state-by-year effects (see Table A.10 in the Appendix).

Overall, our robustness tests reinforce the idea that the availability of credit matters for consumption, independent of the effects of household wealth and income, for those households with negative income shocks and/or limited liquid assets that potentially need or want to borrow to finance their expenditures.

6. 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. While a number of studies have explored this relationship, this study focuses on the more general effects of reduced bank-based credit availability operating through consumer loans rather than primarily emphasizing the effects on homeowners associated with the large declines in home equity. Even though there has been substantial consolidation in the

banking sector, the local branches of banks still play an important role in providing home equity loans and consumer loans to individuals and households.

We examine the relationship between the health of banks with a presence in a given location and the spending behavior of individual households in those locations, paying particular attention to differences in the consumption responses of constrained versus unconstrained households. We identify households that are most likely to be constrained and potentially in need of bank credit to smooth their expenditures using an innovative and multi-dimensional approach that considers households' negative income shocks, having current income below expected future income, and the value of their relative liquid asset holdings. We also construct arguably exogenous measures of bank health based on multi-locational banks with branches in a given location but whose health is unlikely to be affected by local economic conditions. The results show that while a deterioration in bank health results in lower real (food) consumption overall, the effects tend to be strongest for constrained households—especially those with limited liquid assets that can serve as a buffer against negative shocks to income, and are thus most likely to be reliant on bank credit to maintain their desired consumption path.

References:

Abdallah, Chadi, and William Lastrapes. 2012. "Home Equity Lending and Retail Spending: Evidence from a Natural Experiment in Texas." *American Economic Journal: Macroeconomic*s 4(4): 94–125.

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. 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.

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): 459–81. https://library.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?* Kalamazoo, 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.

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.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2006. "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX." 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." 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.

Charles, Kerwin, and Erik Hurst. 2003. "The Correlation of Wealth across Generations." *Journal of Political Economy* 111(6): 1155–1182.

Chava, Sudheer, and Amiyatosh Purnanandam. 2011. "The Effect of Banking Crisis on Bank-Dependent Borrowers." *Journal of Financial Economics* 99(1): 116–135.

Cooper, Daniel. 2013. "House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral." *Review of Economics and Statistics* 95(4): 1183–1197.

Cooper, Daniel, and Karen Dynan. 2016. "Wealth Effects and Macroeconomic Dynamics." *Journal of Economic Surveys* 30(1): 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 14-12. 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 Working Paper Series 2014-17.

Dynan, Karen. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings*

    Papers on Economic Activity, Spring 2012: 299–362.

Gorbachev, Olga and Maria Luengo-Prado. 2019. "The Credit Card Debt Puzzle: The Role of

    Preferences, Credit Risk, and Financial Literacy." Review of Economics and Statistics

    101(2): 294-309.

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.

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.

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.

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.

Johnson, David, Jonathan Parker, and Nicholas Souleles. 2006. "Household Expenditure and the

    Income Tax Rebates of 2001." *American Economic Review* 96(5): 1589–1610.

Kaplan, Greg, Giovanni Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth."

    Brookings Papers on Economic Activity Spring 2014: 77–153.

Kazarosian, Mark. 1997. "Precautionary Savings—A Panel Study." *Review of Economics and*

    Statistics 79(2): 241–247.

Lehnert, Andreas. 2004. "Housing, Consumption, and Credit Constraints." Board of Governors

    of the Federal Reserve System Finance and Economics Discussion Series No. 2004-63.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and

    the Economic Slump." *The Quarterly Journal of Economics* 128(4): 1687–1726.

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.*

Peek, Joe, and Eric S. Rosengren. 2000. "Collateral Damage: Effects of the Japanese Bank

    Crisis on Real Activity in the United States." *The American Economic Review* 90(1):30-45.

Zeldes, Stephen. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation."

    Journal of Political Economy 97(2): 305–346.

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

	Excluding Home Equity	Including Home Equity
	(1)	(2)
ML NPL50	−21.994***	−83.411***
	(8.068)	(11.662)
House Price Growth	0.101***	0.175***
	(0.009)	(0.012)
Bartik Empl. Growth	0.387***	0.549***
	(0.142)	(0.168)
Observations	10974	10974
Adj. R-squared	0.519	0.503

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 identified in our PSID data. The sample that includes home equity is based on the subset of the CCP individuals with mortgage loan data (see Section A.3 for more details). Additional controls include location (MSA) fixed effects and time fixed effects. The dependent variable in the first column is the percent change in real per capita nonmortgage loans. The dependent variable in the last column is the percent change in real per capita nonmortgage loans plus home equity loans. ML NPL50 (median multi-locational nonperforming loan ratio) is the baseline measure of bank health at the MSA level (see Section 2.2). ML NPL50 is divided by 100 for presentation purposes. Robust standard errors are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table 2: Bank Health and Consumption: Initial Results

	(1)	(2)	(3)	(4)	(5)	(6)
	Full Sample			Wealth Sample		
Log Income	0.214***	0.213***	0.213***	0.257***	0.263***	0.263***
	(0.012)	(0.012)	(0.012)	(0.014)	(0.014)	(0.014)
House Price Growth	0.087	0.055	0.057	0.131*	0.161**	0.161**
	(0.054)	(0.053)	(0.053)	(0.069)	(0.069)	(0.068)
Bartik Empl. Growth	0.081	−0.035	−0.005	0.221	0.285	0.275
	(0.694)	(0.686)	(0.689)	(0.878)	(0.878)	(0.878)
NPL50		−1.517**				
		(0.771)				
Inc. Volatility x NPL50		−0.221**				
		(0.091)				
ML NPL50			−1.461***	−1.740***	−1.671***	−1.672***
			(0.527)	(0.553)	(0.547)	(0.547)
Inc. Volatility x ML NPL50			−0.159*	−0.123	−0.121	−0.120
			(0.087)	(0.098)	(0.098)	(0.098)
Total Wealth/Income					1.404***	
					(0.211)	
Total Assets/Income						1.421***
						(0.218)
Total Debt/Income						−1.817***
						(0.539)
Memo: Total BH Effect		−1.712**	−1.602***	−1.842***	−1.772***	−1.771***
		(0.766)	(0.516)	(0.557)	(0.551)	(0.552)
Observations	54109	54109	54109	29431	29431	29431
Adj. R-squared	0.608	0.608	0.608	0.541	0.542	0.542

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

Notes: The dependent variable is log real food consumption. Income is real after-tax household income. NPL50 is the median nonperforming loan ratio (NPL) measure of bank health, and ML NPL50 is the median NPL ratio for multi-locational banks in the local area. Multi-locational banks are the local branches of parent bank holding companies (BHCs) that have less than 5 percent of overall deposits in a given location, but where multi-locational BHCs account for at least a 5 percent share of all local deposits (see Section 2.2 for more details). "Inc. Volatility" is the mean squared deviation of actual household income from predicted household income (See Section 4 for additional details about households' predicted income); and "Wealth/Income" is the ratio of households' assets (bonds, stocks, cash, and IRA holdings along with the net value of any businesses, vehicles, or non-primary-residence real estate) less any nonhousing debt plus housing equity (house value less outstanding mortgage debt), if any, relative to lagged income. Income is multiplied by 100 and bank health measures are divided by 100 for presentation purposes. The "Total BH Effect" combines the direct bank health (BH) effect with the differential effect based on households' income volatility, evaluated at the mean of income volatility for each estimation sample. The wealth sample is restricted to years for which household wealth data are available. Specifications also include controls for age, age squared, family size, family-size squared, and indicators for single-earner and dual-earner households. (Table A.2 reports point estimates for all these additional variables.) Location, 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 3: Bank Health and Consumption: Controlling for Constrained Households

	Neg. Dev.		Neg. Dev. or < 30		Low LA		Low LA or < 30	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Unconst.	Const.	Unconst.	Const.	Unconst.	Const.	Unconst.	Const.
Log Income	0.223***	0.219***	0.229***	0.227***	0.006***	−0.010***	0.261***	0.262***
	(0.014)	(0.015)	(0.015)	(0.015)	(0.001)	(0.002)	(0.015)	(0.015)
ML NPL50	−1.085*	−2.205***	−0.766	−2.353***	−1.626**	−2.150**	−0.978*	−2.253***
	(0.602)	(0.662)	(0.631)	(0.612)	(0.793)	(0.979)	(0.588)	(0.752)
Inc. Volatility x ML NPL50	0.075	−0.155	0.042	−0.112	0.167	−0.430**	0.084	−0.163
	(0.125)	(0.131)	(0.125)	(0.130)	(0.172)	(0.174)	(0.100)	(0.153)
Total Wealth/Income	0.756***	1.693***	0.788***	1.635***	−0.590	2.784***	1.202***	1.572***
	(0.263)	(0.242)	(0.264)	(0.241)	(0.694)	(0.738)	(0.236)	(0.453)
Memo: Total BH Effect	−1.027*	−2.324***	−0.733	−2.438***	−1.493*	−2.494**	−0.911	−2.383***
	(0.587)	(0.676)	(0.616)	(0.625)	(0.786)	(0.972)	(0.588)	(0.751)
Observations	28074		28074		25122		25122	
Adj. R-squared	0.539		0.539		0.525		0.532	

Sources: See Table 2.

Notes: These estimates are based on the wealth sample. The dependent variable is log real food consumption. See Table 2 for additional variable definitions. In the first specification (columns 1 and 2), constrained households are those with a negative income deviation, while in the third specification (columns 5 and 6), constrained households are those with low (below-median) liquid assets relative to income (LA). In the second and fourth specifications, constrained households have either a negative income deviation or low LA, respectively, or the earner(s) in the household are less than 30 years old. All specifications also include controls for local (Bartik) employment growth, local house-price growth, age, age squared, family size, family-size squared, and indicators for single-earner and dual-earner households. Location, 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 and Liquid Asset Holdings

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.226***	0.223***	0.225***	0.228***
	(0.015)	(0.015)	(0.015)	(0.015)
ML NPL50	−2.076***	−2.716***	−0.203	−1.649*
	(0.631)	(0.884)	(0.571)	(0.969)
Inc. Volatility x ML NPL50	0.133	−0.280*	0.022	0.096
	(0.121)	(0.168)	(0.121)	(0.169)
Total Wealth/Income	1.154***	2.782***	1.000***	0.297
	(0.228)	(0.623)	(0.263)	(0.630)
Memo: Total BH Effect	−1.974***	−2.929***	−0.186	−1.576*
	(0.634)	(0.888)	(0.570)	(0.944)
Observations		28074		
Adj. R-squared		0.539		

Sources: See Table 2.

Notes: The dependent variable is real log food consumption. See Table 2 for additional variable definitions. Constrained households have a negative income deviation or all household earners are younger than 30, as discussed in the main text. Constrained and unconstrained households are further grouped by their liquid assets relative to income (LA): "High LA" (above median LA) or "Low LA" (below median LA). LA includes cash holdings (checking and savings accounts), stock holdings, and bond holdings. Specifications also include controls for local (Bartik) employment growth, local house price growth, age, age squared, family size, family size squared, and indicators for single-earner, and dual-earner households. Location, 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: Nonlinear Approach to Bank Health

	(1)	(2)
Log Income	0.264***	0.263***
	(0.014)	(0.014)
ML NPL50 Quint.2	−0.020	−0.021
	(0.013)	(0.013)
ML NPL50 Quint.3	−0.020	−0.019
	(0.016)	(0.015)
ML NPL50 Quint.4	−0.026	−0.020
	(0.018)	(0.019)
ML NPL50 Quint.5	−0.052**	−0.051**
	(0.024)	(0.024)
Inc. Vol. x ML NPL50 Quint.2		0.001
		(0.004)
Inc. Vol. x ML NPL50 Quint.3		0.000
		(0.004)
Inc. Vol. x ML NPL50 Quint.4		−0.008
		(0.005)
Inc. Vol. x ML NPL50 Quint.5		−0.002
		(0.004)
Total Wealth/Income	1.408***	1.412***
	(0.210)	(0.210)
Observations	29431	29431
Adj. R-squared	0.542	0.542

Sources: See Table 2.

Notes: These estimates are based on the wealth sample. The dependent variable is log real food consumption. Specifications allow for possible nonlinear bank health effects by dividing locations based on their position (quintile) in the multi-locational bank health distribution in a given year, with a higher quintile being associated with a higher median nonperforming loan (NPL50) ratio (weaker bank health). Additional control variables are not reported but are the same as in the previous tables. 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: Nonlinear Approach to Bank Health
Controlling for Constrained Households

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.226***	0.224***	0.225***	0.228***
	(0.015)	(0.015)	(0.015)	(0.015)
ML NPL50 Quint.2	−0.014	−0.022	−0.012	−0.026
	(0.019)	(0.031)	(0.021)	(0.027)
ML NPL50 Quint.3	−0.019	−0.038	−0.004	−0.029
	(0.019)	(0.029)	(0.022)	(0.027)
ML NPL50 Quint.4	−0.004	−0.087***	0.018	−0.001
	(0.023)	(0.034)	(0.024)	(0.029)
ML NPL50 Quint.5	−0.062**	−0.085**	0.006	−0.040
	(0.025)	(0.035)	(0.025)	(0.037)
Inc. Vol. x ML NPL50 Quint.2	0.009	−0.008	0.005	0.013
	(0.006)	(0.007)	(0.005)	(0.010)
Inc. Vol. x ML NPL50 Quint.3	−0.001	−0.002	0.004	0.013*
	(0.006)	(0.007)	(0.006)	(0.007)
Inc. Vol. x ML NPL50 Quint.4	0.006	−0.010	0.006	−0.005
	(0.007)	(0.009)	(0.009)	(0.007)
Inc. Vol. x ML NPL50 Quint.5	0.006	−0.002	0.002	0.005
	(0.005)	(0.008)	(0.005)	(0.006)
Total Wealth/Income	1.131***	2.769***	0.971***	0.301
	(0.228)	(0.621)	(0.261)	(0.637)
Observations	28074			
Adj. R-squared	0.540			

Sources: See Table 2.

Notes: The dependent variable is real log food consumption. See Tables 4 and 5 for additional variable definitions. Additional control variables are the same as in the previous tables. 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: Controlling for Constrained, Liquid Assets
Alternative Definition of Constrained

| | | ≥5% Neg. Dev. or < 30 | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.234***	0.229***	0.233***	0.236***
	(0.017)	(0.017)	(0.016)	(0.016)
ML NPL50	−2.606***	−2.691***	−0.202	−1.801*
	(0.773)	(0.958)	(0.608)	(1.027)
Inc. Volatility x ML NPL50	0.189	−0.326*	0.041	0.111
	(0.127)	(0.180)	(0.126)	(0.173)
Total Wealth/Income	1.286***	3.054***	1.041***	0.303
	(0.269)	(0.710)	(0.283)	(0.689)
Memo: Total BH Effect	−2.455***	−2.952***	−0.170	−1.712*
	(0.771)	(0.953)	(0.608)	(1.003)
Observations		25122		
Adj. R-squared		0.533		

Notes: The dependent variable is real log food consumption. Constrained households have a 5 percent or greater (in absolute value) negative income deviation or all earners in a household are younger than 30. Otherwise the specifications are the same as those in Table 4, including the additional control variables. 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

Additional PSID Data Details

- Most PSID variables capture household information as of the survey year. However household income and some other labor market data cover the year prior to the survey year. For example, (total) family income data reported in the 2011 wave are for calendar year 2010.

- 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 are collected in all waves except 1988 and 1989.

- Starting in 1999, the PSID expanded its recorded household expenditure categories to include healthcare, childcare, transportation, school, vehicle, and utilities, in addition to food. Additional categories including home maintenance (upkeep), clothing, recreation, home furnishings, and vacation expenditures were added in 2005.

- Where relevant, we use the quarterly personal consumption expenditures (PCE) deflator (base year 2000) 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.

- Household wealth is the sum of households' housing equity (if any) and their financial assets less any nonhousing debt. Housing equity equals a household's self-reported home value less any outstanding (primary or secondary) mortgage debt. (Housing

equity is zero for renters). Financial assets include the value of a household's bond, stock, cash, and individual retirement account (IRA) holdings, along with the net value of any businesses or farms, vehicles, or non-primary-residence real estate properties owned. Non-housing debt includes a household's student loan debt, credit card debt, personal loans (from banks or relatives), and any other non-real estate related liabilities.

Linking PSID Households with Local Banks

As noted in the text, the PSID unfortunately does not contain information about households' banking relationships. Instead, we link households to banks in their local area using household-level geographic data obtained under a restricted data contract with the PSID designed to protect the confidentiality of the respondents. 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) 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 described in the next subsection), we assign the household to the state's rural (non-MSA) area when linking it with local banking conditions. (New Jersey, Delaware, Rhode Island, and Washington, DC, do not have rural areas, given our location definitions.)

A.1.2 Bank Sample Selection

In constructing bank health, we focus our analysis on FDIC-insured commercial and savings banks headquartered in the 50 US states and the District of Columbia. Since the variables needed for constructing bank health measures are reported at the consolidated bank (parent-bank) level rather than at the local bank branch level, we assign each bank branch the value of the bank

health measure of its parent bank for that period.[1] In addition, rather than equally weighting the health of these bank branches, we use SOD data to construct deposit-weighted measures, since the relative importance of a bank to a given location is arguably tied to its branch's (or branches') share of the total deposits in that location.[2]

We employ the following additional restrictions 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.[3]

- To account for changes in bank composition from quarter to quarter in a given location, for example due to closures or acquisitions—changes not yet reflected in the annual SOD data—we restrict our analysis to branches of parent banks that have Call Report data for a given quarter.

- 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

---

[1] 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; Cortes and Strahan 2015). However, for the purpose of designating multi-locational banks for our bank health measures, we base branch affiliations on the branch's top parent holding company rather than on its parent bank. Internal capital markets also operate within holding companies, and, in any case, we obtain quite similar results using ML bank health measures based on bank-level rather than holding company-level affiliations for the branches.

[2] Weighting based on a branch's loan originations in the local market is preferable but not feasible, given data availability.

[3] 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 rssd9425] does not equal zero.

quarters of observations of voluntarily liquidated banks and the first eight quarters in

which a de novo bank files a Call Report.

Because our primary measure of bank health is based on the distribution of nonperforming loans

and deposits across banks with branches in a given location, we need enough individual banks to

have a representative distribution of banks' financial conditions by location. Consequently, we

exclude MSAs that have fewer than five banks at any time during our sample period.[4]

A.1.3 Constructing Multi-Locational Bank Health

For a given location, we deem a bank branch to be multi-locational if less than 5 percent of

the parent bank holding company's (BHC's) deposits come from that particular location. We

then construct our bank health measures, such as the median nonperforming loan ratio, for the set

of multi-locational BHCs with branches in that location. Note that each BHC is determined to be,

or not be, multi-locational on a location-by-location basis, because a given BHC 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. In addition, in our regression sample we include only the locations in

which the branches of all multi-locational banks account for at least 5 percent of local deposits to

---

[4] 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.

ensure that these multi-locational banks have a potentially meaningful effect on credit availability in that location. [5]

For the purpose of designating bank branches as multi-locational, we use the branch's affiliation with the parent bank holding company rather than the branch's parent bank, given the operation of internal capital markets and the paucity of multi-locational banks based on bank rather than BHC affiliation in the early years of our sample prior to the substantial consolidation of banks after the Riegle-Neal Interstate Banking and Efficiency Act of 1994 which allows banks to branch across state lines.

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 values using the PCE deflator and then define house price growth in a given location as the four-quarter percent change in the location's real house price index. Rural households or households in a metropolitan statistical area (MSA) with no CoreLogic house price data are assigned state-level CoreLogic real house price growth rates, which we construct the same way. We determine house price growth rates based on households' 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

---

[5] We obtain similar results if we broaden our definition of ML banks to include any BHC branches in a location with local deposits that represent 10 percent or less of the parent BHC's overall deposits, if we take a more restrictive approach and include only locations where ML BHCs account for 10 percent or more of local deposits, or if we base the ML designation on bank rather than BHC affiliations of the bank branches (see Section A.4).

percentage 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). [6] 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. [7]

In calculating BEG, we make some adjustments to the QCEW data to address the change over time in industry classifications. 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:

---

[6] See https://www.bls.gov/cew/datatoc.htm for more details on these data.
[7] 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.

- We standardize industries at the county level between the two classification systems into 14 broad industries.[8]

- 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.[9]

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

---

[8] 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.

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

In our empirical specification (equation 1 in the main text), the timing of the data for the 2003 observation of household i interviewed in the first quarter of 2003 (2003 PSID wave) is as follows:

$$c_{2003}^{i,j} = \beta_0 + \beta_1\, y_{2002}^{i,j} + \beta_2\, BH_{2002q1}^{j} + \beta_3 BH_{2002q1}^{j} \cdot V^i + \beta_4\, HPG_{2002q4,2001q4}^{j}$$

$$+ \beta_5\, BEG_{2002q4,2001q4}^{j} + \beta_6\, WY_{2001}^{i,j} + \gamma\, X_{2003}^{i,j} + \delta_{2003} + \eta^i$$

$$+ \epsilon_{2003}^{i,j}. \qquad\qquad (A.1)$$

That is, annualized consumption reported in the first quarter of 2003 by household i living in location j ($c_{2003}^{i,j}$) 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 household wealth relative to income 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 cohabiting households when *both* the head and spouse/partner report being 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-cohabiting) 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 wealth-to-income ratios, as well as those with loan-to-value ratios greater than two. 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 percentage declines in mortgage debt when the initial mortgage balance is less than $25,000.

In addition, we exclude households that report any business income, since the PSID imputes business owners' labor earnings, which tend to be volatile and are often negative. Further, we remove 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.

After applying these sample criteria, we impose a few additional sample 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.[10] We apply the same moving criterion to renters who report that they switch to owning, and we drop observations where households say

---

[10] 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.

they are renters 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 household wealth distribution in each year. We do not adjust the wealth-to-income ratios that we use in equation (1) in the main text for outliers themselves, but instead drop outliers from the numerator and denominator before constructing the ratios. For the location-specific data (NPL50, ML NPL50, HPG, BEG, and any other bank health measures), we winsorize the data at the 1^{st} and 99^{th} 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.

A.2 MSA-level Loan Growth Analysis

A.2.1 Data

We calculate real per capita non-mortgage loans in MSAs using the Federal Reserve Bank of New York Consumer Credit Panel (CCP) provided by Equifax. The CCP is a 5 percent random sample of all individuals in the United States with credit records in 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 CPP 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 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.[11] Individuals with invalid zip codes are dropped from the calculations.[12] 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.[13] Loans marked as "joint" are divided in half to avoid potential double counting.[14] 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

---

[11] 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.

[12] 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).

[13] 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.

[14] 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."

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.

A.2.2 Estimation Equation

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

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

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 measure of bank health (e.g., ML NPL50) as of *t-1* in MSA *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.

A.3 Calculating Age-Earnings Profiles and Households' Income Deviations

A.3.1 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:

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

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 (A.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 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.[15] 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, or being in school), we have a number of procedures for cleaning the data to generate more accurate estimates of $\beta_1^{\omega,s}$ and $\beta_2^{\omega,s}$ that we describe in more detail below in Section A.3.3.

We use our estimates $\hat{\beta}_1^{\omega,s}$ and $\hat{\beta}_2^{\omega,s}$ 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, the intercepts may differ due to the careers (industries) the individuals choose. We construct the individual-specific age-earnings profile intercept, α_i, for each observation for a given individual as:

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

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 observation (age and income combination) for a given individual

---

[15] We assign individuals to the education group consistent with their highest observed education level.

will potentially yield a different value 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.

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

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 person.[16]

A.3.2 Estimating Households' Income Deviations

Our variable of interest for determining likely income-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:

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

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

A.3.3 Estimating Households' Predicted Income: Additional Details[17]

---

[16] 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 observations. However, for those individuals who have at least five observations with labor income so that we calculate an age-earnings profile, we include observations where the individual is unemployed for the purpose of calculating the individual's income deviation.

[17] There are some slight nuances to the procedures described in this subsection, and additional information is available from the authors upon request.

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.

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.[18]

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

---

[18] Note 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.

structure of the PSID, the age of an individual cannot realistically decrease between years.

Therefore, when we see a non-temporary reported 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

tempoarily declines in one year and then reverts to its original path, we adjust only the 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.[19] We identify these individuals

in the PSID as follows:

---

[19] 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. In addition, business owners' labor income is imputed, and often negative.

- 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.[20] 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

---

[20] This is likely a situation in which unemployment is reported or recorded incorrectly.

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.[21]

- • Unemployed Workers: We treat individuals who report that they are unemployed somewhat differently.[22] 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.4 Robustness Tests

A.4.1 Alternative Measures of Bank Health

Alternative Definitions using Call Report Data

The median multi-locational (ML) nonperforming loan (NPL) ratio is not the only way to capture banks' relative health using bank balance sheet data. Given our finding that it is the banks in the highest (worst) quintile of the bank health distribution that appear to be driving our overall results, alternative balance-sheet-based measures of bank health that also emphasize the

---

[21] 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.
[22] We classify individuals as unemployed if they report one week (or more) of unemployment in a given period.

role of the health of banks in the tail of the distribution may be informative. Thus, we consider measures from the lower end of the ML NPL distribution, such as the 25th or 10th percentiles, rather than ML NPL50. However, these measures produce results that are very similar to our main findings in Table 4 in the text.[23] In part, the similarity of the results occurs because the median, 25th percentile, and 10th percentile of the ML NPL distribution are highly correlated. This makes sense, insofar as when a bank in the middle of the distribution is under financial stress—for example, due to current and prospective losses associated with a sharp increase in its nonperforming loans—then many banks in the location are probably also stressed. Credit availability in that location therefore likely shrinks as the bank at the median along with other banks in even worse shape, such as those at the 25th and 10th percentiles of the distribution, curtail their lending to improve their balance sheets.

We also consider an alternative bank health measure based on the ML nonperforming *real estate* loan (RENPL) ratio, given that both the recent financial crisis as well as the banking crisis in the late 1980s and early 1990s were characterized by severe real estate loan problems. Indeed, even though our baseline (total) NPL measure includes nonperforming real estate loans, banks' RENPL ratio, which emphasizes a particular type of loan on banks' balance sheets, might be a better gauge of the state banks' health. Table A.4 shows the same specifications as in Table 4 in the main text, but uses variation in the median ML RENPL ratio (ML RENPL50) across time and locations to examine the effect of bank health on consumption. The estimates are very similar to, if not at touch stronger than, our baseline findings. Indeed, the consumption of constrained households with low (below median) liquid assets relative to income (LA) declines the most in response to a deterioration in banks' RENPL ratios, while the consumption of

---

[23] These results are not shown, since they are very similar to our main findings, but are available from the authors upon request.

constrained households with high (above median) LA also falls, but by only about two-thirds as much. The consumption of unconstrained households decreases as well, but to a lesser extent. Like we discuss in the main text, these results could be capturing the fact that some households may have been subjected to unexpected expenditure shocks or may appear wealthy or unconstrained but behave otherwise to better position themselves should a shock occur in the future when bank credit may be less available.

In addition, our approach for measuring ML bank health requires assumptions about what makes a given banking organization multi-locational. Recall that we define a ML bank as one where deposits at the banks' branches in a given location are 5 percent or less of the parent BHC's overall deposits. We find very similar results, however, if instead, we broaden our definition of ML banks to include any BHC with branches in a location where local deposits account for 10 percent or less of its overall deposits (see Table A.5).[24] Remember as well that our sample selection criteria limits our analysis to locations where the branches of all ML banks in the location account for at least 5 percent of the local deposits. Our results are also very similar, however, if we take a more restrictive approach and include only locations where ML banks account for 10 percent or more of the deposits (see Table A.6). Finally, designating ML branches using bank-level affiliations rather than BHC affiliations also yields very similar results to our baseline findings (not shown). Overall, our results are robust to alternative approaches to measuring ML banks' health based on balance sheet variables.

Bank Supervisory Data

---

[24] We use the 5 percent deposit share cutoff as a baseline for identifying ML banks because the bank health measure is arguably more exogenous with respect to local conditions for banks with a smaller share of their overall deposits in a given location.

As a further robustness check, we measure bank health based on confidential 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.[25] Because the CAMELS rating values are integers rather than a continuous variable and have limited range, we use a weighted average rather than the median value as with the nonperforming loan ratio. We follow the same procedure as described in the main text and above for determining ML bank health based on the weighted average CAMELS rating data.

When we use the CAMELS data rather than bank balance sheet data to measure bank health, the results are very similar to our baseline results (see Table A.7). In particular, the point estimate of the sensitivity of food consumption to changes in ML bank health, as captured by the supervisory data, for constrained households with low LA is almost double the sensitivity for constrained households with high LA. (Note that the magnitude of the bank health effects with the CAMELS data are different than with the Call Report data because they are based on an entirely different scale.) In comparison, bank health based on the CAMELS ratings has little effect on the consumption of unconstrained households—especially unconstrained households with high LA.

---

[25] 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 (de novo) bank have already been omitted from our bank health measures.

Obtaining results very similar to our main findings in the text using a completely different approach to measuring bank health is reassuring. It suggests that our findings are not being driven by our decision to focus on banks' balance sheet data to capture their relative health.

A.4.2 Alternative Measures of Consumption

Food may not be the ideal expenditure measure to study household consumption. However, food consumption's redeeming feature in the PSID is that the data are available for a much longer time period than are broader consumer expenditure data, which importantly allows us to include data from the banking crisis in the late 1980s and early 1990s in our analysis. Moreover, the food data include a discretionary component (food away from home) that likely fluctuates with household income and other factors.

Indeed, we find broadly similar results if we augment the food expenditure data with additional, likely discretionary, expenditure data that are available in the PSID starting in 2005 (see Table A.8). These data include household expenditures on home furnishings, recreation, vacation, school costs, and clothing—both more discretionary and potentially bigger-ticket spending categories where households may be more likely to turn to bank credit for financing (see Section A.1.1 for further details about these additional data). Despite the shorter sample, the results in Table A.8 are broadly similar to our main findings. In particular, the effect of bank health on this broader (discretionary) consumption measure is strongest and most precisely estimated for constrained households with low LA, while the bank health effects for other household types are smaller (in absolute value) and not statistically significant. It is likely that the much smaller sample size that is split across four household groups contributes to the difficultly of obtaining statistically significant effects. However, since the sample period for

these estimates covers the Great Recession and its aftermath, it still includes an episode when bank health was severely impaired, and thus we are able to identify the effect of a deterioration in bank health on consumption.

An alternative approach, which allows for estimates over a longer sample horizon in the PSID using a broader measure of consumption, is to impute a nondurable 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 given household characteristics and the available data on food expenditures. Qualitatively, the estimated bank health effects are similar to our baseline findings—bank health continues to matter most for constrained households' consumption (see Table A.9).

A.4.3 Additional Robustness Checks

An additional potential concern with our results is that our location-specific bank health controls are simply picking up time-varying differences in macroeconomic conditions across locations that are not well captured by our controls for local economic conditions. That is, household consumption may be stronger over time in locations with healthier banks compared with locations with less healthy banks due to unobserved economic factors unrelated to bank health.

While we do not believe such unobserved factors are driving our results, we address this concern by re-estimating our results after including state-by-year effects. Including these additional fixed effects, as expected, has little effect on the patterns of our estimated bank health effects (see Table A.10). In particular, the greatest sensitivity of consumption to bank health

occurs for constrained households with low LA, and constrained households have greater

sensitivity to bank health than do unconstrained households. As a result, it does not appear that

our findings are driven by unobserved differences across locations that vary over time.[26]

---

[26] Ideally, we would include MSA-by-year effects to check the robustness of our results. However, since we have some states with only one or two MSAs, including such effects soaks up too much of the variation in our location-specific, time varying bank health measure.

Table A.1: Summary Statistics

Variable	Full Sample						Wealth Sample					
	Mean	Median	Std. Dev	Min.	Max.	N	Mean	Median	Std. Dev	Min.	Max.	N
Food Consumption [$]	5,541	5,126	3,364	11	51,566	54,109	5,880	5,422	3,296	74	22,151	29,431
Imputed Nondurable Consumption [$]	14,126	13,521	7,104	48	78,577	49,731	16,155	15,593	6,961	294	48,120	27,300
Alt PSID Consumption [$]	11,278	9,217	9,500	74	323,939	20,251	11,226	9,179	9,511	74	323,939	18,669
Income [$]	49,048	43,086	30,299	748	300,554	54,109	52,942	46,674	32,228	748	300,554	29,431
NPL50	1	.69	.9	.13	5.1	54,109	1.1	.74	.98	.13	5.1	29,431
ML NPL50	1.3	.8	1.2	0	9.3	54,109	1.4	.9	1.3	.0074	9.3	29,431
CAMELS	1.9	1.9	.42	1	4.4	68,535	1.9	1.9	.34	1	4.4	35,164
House Price Growth [%]	.96	1.1	6.4	-29	30	54,109	.51	1	7.4	-29	30	29,431
Bartik Empl. Growth [%]	1.4	1.8	1.9	-5.9	5.1	54,109	1.1	1.6	2.1	-6	5.1	29,431
Age	42	41	11	17	91	54,109	44	43	11	19	91	29,431
Family Size	3	3	1.5	1	14	54,109	2.9	3	1.5	1	14	29,431
Marital Status: Single	.34	0	.47	0	1	54,109	.34	0	.47	0	1	29,431
Single Earner	.18	0	.39	0	1	54,109	.19	0	.39	0	1	29,431
Dual Earner	.48	0	.5	0	1	54,109	.47	0	.5	0	1	29,431
Const. HHs (Neg. Dev. or < 30) [%]	53	49	7.1	44	62	51,965	48	47	5.1	44	61	28,074
Const. HHs (≥5% Neg. Dev. or < 30) [%]	47	42	6.8	38	57	46,302	43	42	4.7	38	55	25,122
Homeowners [%]	64	65	2.5	58	67	54,109	64	64	.94	60	65	29,431
Total Wealth/Income							1.5	.66	2.5	-4.4	34	29,431
Liquid Assets/Income							.32	.041	.93	0	20	29,431
Liquid Assets [$]							23,188	1,899	81,135	0	2,345,770	29,431

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

Notes: Imputed nondurable consumption is based on the approach in Blundell, Pistaferri, and Preston (2006). The alternative consumption measure consists of food, clothing, school, recreation, vacations, and furnishings. Income is after-tax real household income. Dollar measures are in constant 2000 dollars using the PCE price deflator. NPL50 is the median nonperforming loan ratio (NPL) measure of bank health (BH). ML NPL50 is the median NPL measure for multi-locational (ML) banks, which are the bank branches in a local area that account for less than 5 percent of their parent bank holding company's (BHC) overall deposits, but where ML BHC deposits account for at least a 5 percent share of all local deposits (see Section 2.2 for more details). CAMELS is the deposit-weighted average of bank supervisors' confidential bank health ratings (integer scale of 1 to 5). 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). See main text for details on how constrained households are defined. Wealth/Income is the ratio of households' assets (bonds, stocks, cash, and IRA holdings along with the net value of any businesses, vehicles, or non–primary-residence real estate) less any nonhousing debt plus housing equity (house value less outstanding mortgage debt). Liquid Assets/Income is the ratio of liquid assets (cash holdings [checking and savings accounts], stock holdings, and bond holdings) to lagged income.

Table A.2: Correlations of Non-Demographic Regressors
Full Sample

	Log Income	House Price Growth	Bartik Empl. Growth	NPL50	Inc. Volatility x NPL50	ML NPL50	Inc. Volatility x ML NPL50
Log Income	1						
House Price Growth	-0.0122**	1					
Bartik Empl. Growth	-0.0557***	0.450***	1				
NPL50	0.0325***	-0.256***	-0.0991***	1			
Inc. Volatility x NPL50	-0.0573***	-0.0335***	-0.0144***	0.163***	1		
ML NPL50	0.0493***	-0.218***	-0.0865***	0.799***	0.122***	1	
Inc. Volatility x ML NPL50	-0.0533***	-0.0312***	-0.0121**	0.143***	0.933***	0.159***	1

Wealth Sample

	Log Income	House Price Growth	Bartik Empl. Growth	NPL50	Inc. Volatility x NPL50	ML NPL50	Inc. Volatility x ML NPL50	Total Wealth /Income	Total Assets /Income	Total Debt /Income
Log Income	1									
House Price Growth	-0.00502	1								
Bartik Empl. Growth	-0.0462***	0.491***	1							
NPL50	0.0190**	-0.235***	-0.0156**	1						
Inc. Volatility x NPL50	-0.0579***	-0.0372***	-0.00481	0.176***	1					
ML NPL50	0.0342***	-0.187***	0.00762	0.835***	0.132***	1				
Inc. Volatility x ML NPL50	-0.0561***	-0.0319***	-0.000639	0.158***	0.947***	0.160***	1			
Total Wealth/Income	0.234***	-0.0397***	-0.0384***	-0.00281	-0.0127*	0.000828	-0.0150*	1		
Total Assets/Income	0.302***	-0.0603***	-0.0618***	0.0331***	-0.00986	0.0307***	-0.0126*	0.916***	1	
Total Debt/Income	0.245***	-0.0570***	-0.0669***	0.0922***	0.00526	0.0777***	0.00389	0.0475***	0.431***	1

Sources: See Table A.1.

Notes: See Table A.1 for definitions.

Table A.3: Bank Health and Consumption: Initial Results
All Regressors

	(1)	(2)	(3)	(4)	(5)	(6)
	Full Sample			Wealth Sample		
Log Income	0.214***	0.213***	0.213***	0.257***	0.263***	0.263***
	(0.012)	(0.012)	(0.012)	(0.014)	(0.014)	(0.014)
House Price Growth	0.087	0.055	0.057	0.131*	0.161**	0.161**
	(0.054)	(0.053)	(0.053)	(0.069)	(0.069)	(0.068)
Bartik Empl. Growth	0.081	−0.035	−0.005	0.221	0.285	0.275
	(0.694)	(0.686)	(0.689)	(0.878)	(0.878)	(0.878)
NPL50		−1.517**				
		(0.771)				
Inc. Volatility x NPL50		−0.221**				
		(0.091)				
ML NPL50			−1.461***	−1.740***	−1.671***	−1.672***
			(0.527)	(0.553)	(0.547)	(0.547)
Inc. Volatility x ML NPL50			−0.159*	−0.123	−0.121	−0.120
			(0.087)	(0.098)	(0.098)	(0.098)
Total Wealth/Income					1.404***	
					(0.211)	
Total Assets/Income						1.421***
						(0.218)
Total Debt/Income						−1.817***
						(0.539)
$Age^2/1000$	−0.287***	−0.288***	−0.288***	−0.375***	−0.391***	−0.394***
	(0.040)	(0.040)	(0.040)	(0.049)	(0.049)	(0.049)
Family Size	0.199***	0.199***	0.199***	0.180***	0.178***	0.178***
	(0.017)	(0.017)	(0.017)	(0.018)	(0.018)	(0.018)
Family Size Squared	−0.016***	−0.016***	−0.016***	−0.016***	−0.016***	−0.016***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Single Earner	0.016	0.016	0.016	0.020	0.014	0.015
	(0.025)	(0.025)	(0.025)	(0.030)	(0.030)	(0.030)
Dual Earner	0.023	0.023	0.023	0.033	0.031	0.032
	(0.023)	(0.023)	(0.023)	(0.028)	(0.028)	(0.028)
Memo: Total BH Effect		−1.712**	−1.602***	−1.842***	−1.772***	−1.771***
		(0.766)	(0.516)	(0.557)	(0.551)	(0.552)
Observations	54109	54109	54109	29431	29431	29431
Adj. R-squared	0.608	0.608	0.608	0.541	0.542	0.542

Sources: See Table A.1.

Notes: The dependent variable is log real food consumption. Table is equivalent to Table 2 in the main text, except that it reports point estimates for all explanatory variables except the fixed effects. "Single Earner" is a dummy variable that takes a value of 1 if there are two (married or cohabitating) adults in the household but only one earner. "Dual Earner" is a dummy variable that takes a value of 1 if there are two (married or cohabitating) adults and both are wage earners (the excluded category is single (one adult) households). Other variable definitions are the same as in Table 2 in the main text. Location, 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 A.4: Bank Health and Consumption: Bank Health Based on Real Estate Nonperforming Loans

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.226***	0.222***	0.226***	0.228***
	(0.015)	(0.015)	(0.015)	(0.015)
ML RENPL50	−2.210***	−3.190***	−0.380	−1.812*
	(0.763)	(0.958)	(0.676)	(1.069)
Inc. Volatility x ML RENPL50	0.144	−0.352*	0.011	0.080
	(0.129)	(0.183)	(0.141)	(0.186)
Total Wealth/Income	1.144***	2.781***	0.989***	0.290
	(0.228)	(0.623)	(0.262)	(0.630)
Memo: Total BH Effect	−2.101***	−3.458***	−0.372	−1.751*
	(0.768)	(0.960)	(0.672)	(1.042)
Observations		28075		
Adj. R-squared		0.539		

Sources: See Table A.1.

Notes: The dependent variable is log real food consumption. RENPL50 is the median real estate nonperforming loan ratio in a given location (alternative measure of bank health). Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.5: Bank Health and Consumption: Alternative Def. ML Banks
Parent Banks with 10% or Less of Overall Deposits in Location

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.228***	0.225***	0.228***	0.231***
	(0.016)	(0.016)	(0.016)	(0.016)
ML NPL50	−2.248***	−3.271***	−0.515	−2.035*
	(0.757)	(0.951)	(0.643)	(1.057)
Inc. Volatility x ML NPL50	0.128	−0.281	0.035	0.140
	(0.115)	(0.173)	(0.126)	(0.171)
Total Wealth/Income	1.180***	2.623***	0.878***	0.145
	(0.243)	(0.660)	(0.297)	(0.655)
Memo: Total BH Effect	−2.148***	−3.490***	−0.488	−1.925*
	(0.756)	(0.952)	(0.640)	(1.025)
Observations		25792		
Adj. R-squared		0.541		

Sources: See Table A.1.

Notes: The dependent variable is log real food consumption. In these specifications, multi-locational (ML) banks are defined as the bank branches in a local area that account for less than *10* percent of their parent BHC's overall deposits, but where ML BHC deposits account for at least a 5 percent share of all local deposits. Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.6: Bank Health and Consumption: Alternative Definition ML Banks
ML Banks Have 10% Share (or More) of Local Deposits

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.226***	0.222***	0.226***	0.229***
	(0.015)	(0.015)	(0.015)	(0.015)
ML NPL50	−2.559***	−2.950***	−0.613	−2.424**
	(0.710)	(0.982)	(0.646)	(1.047)
Inc. Volatility x ML NPL50	0.100	−0.237	0.003	0.052
	(0.118)	(0.182)	(0.122)	(0.168)
Total Wealth/Income	1.144***	2.776***	0.988***	0.256
	(0.227)	(0.627)	(0.260)	(0.626)
Memo: Total BH Effect	−2.483***	−3.131***	−0.610	−2.384**
	(0.707)	(1.002)	(0.634)	(1.029)
Observations		28074		
Adj. R-squared		0.539		

Sources: See Table A.1.

Notes: The dependent variable is log real food consumption. In these specifications, multi-locational (ML) banks are defined as the bank branches in a local area that account for less than 5 percent of their parent BHC's overall deposits, but where ML BHC deposits account for at least a *10* percent share of all local deposits. Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.7: Bank Health and Consumption: Alternative (Supervisory-based) Bank Health Measure

	(1)	(2)	(3)	(4)
	Neg. Dev. or < 30			
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.216***	0.220***	0.206***	0.212***
	(0.014)	(0.014)	(0.013)	(0.014)
CAMELS	−0.067***	−0.114***	−0.001	−0.027
	(0.021)	(0.025)	(0.019)	(0.029)
Inc. Volatility x CAMELS	−0.004	−0.009*	−0.004	−0.008
	(0.005)	(0.005)	(0.004)	(0.006)
Total Wealth/Income	1.114***	3.011***	1.083***	−0.111
	(0.212)	(0.568)	(0.224)	(0.620)
Memo: Total BH Effect	−0.070***	−0.121***	−0.004	−0.033
	(0.021)	(0.026)	(0.019)	(0.029)
Observations	33453			
Adj. R-squared	0.542			

Sources: See Table 2 in the main text.

Notes: The dependent variable is real log food consumption. CAMELS is the deposit-weighted average of bank supervisors' confidential bank health ratings (integer scale of 1 to 5). Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.8: Bank Health and Consumption: Alternative Consumption Measure, 2005-2015

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.265***	0.265***	0.267***	0.268***
	(0.025)	(0.025)	(0.024)	(0.025)
ML NPL50	−1.430	−2.661**	−1.399	−1.358
	(1.038)	(1.070)	(0.986)	(1.193)
Inc. Volatility x ML NPL50	−0.158	−0.069	−0.018	0.092
	(0.142)	(0.157)	(0.136)	(0.226)
Total Wealth/Income	0.335	1.518*	0.910*	0.059
	(0.431)	(0.812)	(0.476)	(0.817)
Memo: Total BH Effect	−1.555	−2.715**	−1.413	−1.286
	(1.043)	(1.065)	(0.988)	(1.202)
Observations			17713	
Adj. R-squared			0.640	

Sources: See Table 2 in the main text.

Notes: The dependent variable is a more comprehensive (and discretionary) measure of households' real consumption (sum of food, clothing, school, recreation, vacations, and home furnishing expenditures). These additional data are available only from 2005 onward. Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.9: Bank Health and Consumption: Imputed Nondurables Consumption

| | Neg. Dev. or < 30 | | | |
| | (1) | (2) | (3) | (4) |
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.126***	0.122***	0.125***	0.126***
	(0.011)	(0.011)	(0.011)	(0.011)
ML NPL50	−1.084**	−0.900*	−0.172	−0.965*
	(0.447)	(0.543)	(0.427)	(0.551)
Inc. Volatility x ML NPL50	0.044	−0.136	0.011	0.041
	(0.076)	(0.102)	(0.059)	(0.102)
Total Wealth/Income	0.333**	1.309***	0.416**	0.230
	(0.156)	(0.367)	(0.172)	(0.387)
Memo: Total BH Effect	−1.050**	−1.003*	−0.163	−0.934*
	(0.445)	(0.541)	(0.427)	(0.544)
Observations		25928		
Adj. R-squared		0.622		

Sources: See Table 2 in the main text.

Notes: The dependent variable is imputed nondurable household consumption based on the approach in Blundell, Pistaferri, and Preston (2006). Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

Table A.10: Bank Health and Consumption: State x Year Fixed Effects

	Neg. Dev. or < 30			
	(1)	(2)	(3)	(4)
	Const., High LA	Const., Low LA	Unconst., High LA	Unconst., Low LA
Log Income	0.224***	0.220***	0.222***	0.225***
	(0.015)	(0.015)	(0.015)	(0.015)
ML NPL50	−2.246**	−2.578**	−0.133	−1.357
	(0.935)	(1.035)	(0.821)	(1.102)
Inc. Volatility x ML NPL50	0.149	−0.258	0.061	0.114
	(0.121)	(0.172)	(0.120)	(0.181)
Total Wealth/Income	1.052***	2.569***	0.920***	0.077
	(0.236)	(0.647)	(0.272)	(0.653)
Memo: Total BH Effect	−2.132**	−2.775***	−0.086	−1.270
	(0.940)	(1.037)	(0.812)	(1.094)
Observations		28055		
Adj. R-squared		0.539		

Sources: See Table 2 in the main text.

Notes: The dependent variable is real log food consumption. Regressions include state x year fixed effects as discussed in the main text. Otherwise the specifications are the same as those in Table 4 in the main text, including the additional control variables. Robust standard errors clustered by location are in parentheses: *, **, and *** indicate significance at the 10 percent, 5 percent, and 1 percent levels, respectively.