

w o r k i n g
p a p e r

15 07

**The Rise and Fall of Consumption
in the '00s**

Yuliya Demyanyk, María José Luengo-Prado,
Dmytro Hryshko, and Bent E. Sørensen



FEDERAL RESERVE BANK OF CLEVELAND

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Working papers are available on the Cleveland Fed's website at:

www.clevelandfed.org.

The Rise and Fall of Consumption in the '00s

Yuliya Demyanyk, María José Luengo-Prado,
Dmytro Hryshko, and Bent E. Sørensen

U.S. consumption has gone through steep ups and downs since the turn of the millennium, but the causes of these fluctuations are still imperfectly identified. We quantify the relative impact on consumption growth of income, unemployment, house prices, credit scores, debt, expectations, foreclosures, inequality, and refinancings for four subperiods: the “dot-com recession” (2001-2003), the “subprime boom” (2004-2006), the Great Recession (2007-2009), and the “tepid recovery” (2010-2012). We document that the explanatory power of different factors varies by subperiods, implying that a successful modeling of this decade needs to allow for multiple causal determinants of consumption.

JEL Codes: E21, E24.

Keywords: consumption growth, wealth effects, income inequality, debt overhang, consumer credit, consumer expectations, foreclosures, cash-out refinancing, dot-com recession, subprime boom, Great Recession, tepid recovery.

Suggested citation: Demyanyk, Yuliya, María José Luengo-Prado, Dmytro Hryshko, and Bent E. Sørensen, 2015. “The Rise and Fall of Consumption in the ‘00s,” Federal Reserve Bank of Cleveland, working paper no. 15-07.

Yuliya Demyanyk is at the Federal Reserve Bank of Cleveland (yuliya.demyanyk@clev.frb.org); Dmytro Hryshko is at the University of Alberta (dmytro.hryshko@ualberta.ca), María José Luengo-Prado is at the Federal Reserve Bank of Boston (maria.luengo-prado@bos.frb.org), and Bent E. Sørensen is at the University of Houston and CEPR (besorensen@uh.edu). The authors thank Danny Kolliner for outstanding research assistance. The views expressed are those of the authors and do not necessarily reflect the official positions of the Federal Reserve Bank of Boston, the Federal Reserve Bank of Cleveland, or the Federal Reserve System.

1 Introduction

Private consumption is the largest component of Gross Domestic Product (GDP), constituting over 68 percent in 2014, and its oscillations in the first decade of the new millennium provided much fuel for business cycle fluctuations. In spite of recent influential research which has pinpointed partial explanations of the ups and downs of consumption during the subprime boom and the Great Recession—in particular, the role of housing wealth, subprime lending, and debt overhang—the role of other potential drivers of consumption growth remains mostly unexplored.

The decade was unusually volatile. For example, the start of the millennium witnessed dramatic changes in gross housing wealth, which fell from \$20.7 trillion in 2007 to \$16.4 trillion in 2011 before recovering to \$17.5 trillion in 2012. When house prices fell, owners who fell behind with their mortgage payments were not able to sell their homes for more than they owed and foreclosures ballooned from fewer than 800,000 in 2006 to 2.4 million in 2009. Personal real debt per capita increased steeply from \$31,000 in 2000 to \$56,000 in 2008, when it started gradually declining, reaching \$47,000 in 2012. Confidence eroded steeply from an index value of 106 in the third quarter of 2007 to an exceptionally pessimistic 30 in the first quarter of 2009, before gradually climbing back to 80 at the end of 2012. Unemployment shot up from 5 percent in the fourth quarter of 2008 to 8.2 percent just a year later, peaking at 9.9 percent at the end of 2009 before slowly recovering, ending 2012 at 7.8 percent. Stock market investors lost a staggering \$4+ trillion as the capitalization of the S&P500 index dropped from about \$13 trillion at the end of 2007 to about \$7.8 trillion at the end of 2008. The stock market, however, recovered almost all lost ground by the end of 2012.¹

¹Gross housing wealth is from the Federal Reserve Board’s annual statistical release; real debt per capita is calculated by the authors by aggregating individual-level total debt reported by the Equifax Consumer Credit Panel maintained by the Federal Reserve Bank of New York; population data are from the Census Bureau; foreclosures are from the Mortgage Bankers Association; the Consumer Confidence index is from the Conference Board; the unemployment rate is from the Bureau of Labor Statistics; and stock market capitalization is from Standard and Poor’s.

This paper provides empirical evidence on the relative impact of these factors on consumption growth using U.S. county-level data. We document how consumption growth correlated with income, unemployment, debt, income inequality, consumer expectations, housing wealth, access to credit, cash-out refinancings, and foreclosures. We perform regressions of three-year consumption growth rates on economic variables at the county level, with some variables measured at a higher level of geographic aggregation.² We consider separately four subperiods: the “dot-com recession” (2001-2003), the “subprime boom” (2004-2006), the Great Recession (2007-2009), and the “tepid recovery” (2010-2012).³ We find significant impacts of income and unemployment in all subperiods, while other variables were significant in the subprime boom, the Great Recession, and the tepid recovery. Our results are potentially important for fiscal and monetary policy; for example, if debt overhang explains a large fraction of the variation in consumption, an interest rate policy that lowers debt service may be a powerful stabilizer, but if unemployment is more important for consumption growth, fiscal policy—for example, in the form of increased public purchases—may be more effective. We do not formulate a full model and therefore we do not attempt to suggest optimal policy; however, we provide a set of important stylized facts that a policy-relevant model must match.

Our regressions, being cross-sectional, measure deviations from perfect risk sharing—the case in which consumption growth rates are the same in all counties.⁴ If deviations from perfect risk sharing are economically significant, policymakers face more complicated challenges relative to a world that can be approximated by a representative-agent model. For example, in the subprime

²We use three-year growth rates in order to avoid estimating short-term adjustment using a number of lags. Our interest here is on the medium-run adjustment, and detailed modeling of short-term adjustment would only obscure the results. The choice of geographical observation units is dictated by data availability.

³The label 2001-2003 refers to consumption growth from the year 2000 to the year 2003 approximated by the difference between annual log-consumption in 2003 and annual log-consumption in 2000. The same convention applies for the other subperiods.

⁴Atif Mian, Kamalesh Rao & Amir Sufi (2013) elaborate on this issue in a related paper which finds highly significant (three-year) consumption elasticities with respect to housing wealth during the Great Recession.

boom house prices surged in Arizona and Nevada but not in Texas, raising the question whether an aggregate policy designed to benefit states with surging house prices would also benefit other states—a problem that is more clearly seen in Europe where northern Europe and southern Europe have been diverging since the onset of the Great Recession. Our results are quite heterogeneous across subperiods, which can be seen as an illustration of the Lucas critique, which points out that policy guidance cannot be independent of modeling in times of large underlying structural changes in the economy, when correlations between policy-relevant variables are unstable. Our results are suggestive of the factors that impacted the aggregate U.S. economy in 2000s, although we do not explicitly model aggregate consumption because of the limited degrees of freedom that aggregate data can provide.

We find that income and employment are robust predictors of consumption throughout our sample with the expected signs. Debt overhang is significant in all periods, albeit with a different sign in the dot-com recession. In most periods, consumer confidence and foreclosures are important, while the share of subprime borrowers is important only during the subprime boom, when it is very significant. Growth of cash-out refinancings played an important role during the tepid recovery.

There is large variation in consumption growth rates among counties with different incomes, debt, unemployment, etc. For example, during the dot-com recession, counties with nondeclining income had virtually unchanged consumption on average, while counties with the largest declines in income saw consumption decline by almost 2 percent.⁵ The difference between counties with the highest and lowest increase in unemployment is striking: counties with the lowest increase had a much lower decline of consumption (−0.1 percent) than counties with the largest increase (−5 percent consumption growth on average). Consumption grew by 1.5 percent in counties with the highest growth in housing wealth, whereas consumption declined by 2 percent in counties with the slowest growth in housing wealth. A similar pattern is observed for

⁵The comparison is between counties with income growth above the top 10 and below the bottom 10 percentiles.

consumer confidence—counties where confidence increased the most expanded their consumption, while counties where confidence decreased the most contracted their consumption. During the Great Recession, when all counties had a severe decline in consumption, counties in the top decile of the income distribution, which had positive income growth, had a 30 percent smaller decline in consumption growth compared to counties in the bottom decile of the income distribution. During the tepid recovery, counties with positive growth in cash-out refinancings had on average 3.4 percentage points larger consumption growth than counties at the bottom decile of cash-out-refinance growth.

Analyzing all variables simultaneously, we are able to measure their partial contributions to consumption growth in each of the subperiods. Based on a partial R^2 analysis, we find that unemployment and debt overhang have the highest explanatory power throughout the sample, while the impact of other factors varies by subperiods. The share of income received by the top 10 percentile was important in the dot-com recession while consumer confidence was important in the subprime boom. During the Great Recession, no variables were important beyond unemployment and debt overhang, while income growth, foreclosures, and cash-out refinancings were relevant during the tepid recovery. Further, inequality, housing wealth, and access to credit were important in some of the four subperiods.

The paper is laid out as follows: Section 2 relates our findings to the existing body of literature and Section 3 outlines the relevant theory of consumption. Section 4 describes our data, and Section 5 describes the economy in the four periods we study. Section 6 outlines our empirical specification and describes the results, and Section 7 concludes.

2 Previously established patterns

The economy of the 2000s was dominated by the boom and bust in housing and a boom and bust of subprime mortgages (Yuliya Demyanyk & Otto Van Hemert 2011). Easy credit followed by tight credit, and (housing) wealth gains followed by wealth losses are prime candidates for explaining consumption

patterns. Mian, Rao & Sufi (2013) estimate the consumption elasticity with respect to housing net worth and show that ZIP codes which experienced large wealth losses significantly curtailed consumption. Matteo Iacoviello (2011) discusses the literature on housing wealth effects more broadly and points out that regressions of aggregate consumption on housing wealth may find correlations that are driven by left-out variables. Studies using micro data estimate an elasticity of around 10 percent, although the magnitude is likely to depend on the ease with which homeowners can borrow against housing wealth. Non-housing wealth effects are often found to be smaller. Atif Mian & Amir Sufi (2009) zoom in on the easing of credit conditions associated with subprime lending peaking in the years 2004-2006 and the subsequent bust and debt overhang. They show that mortgage defaults in the Great Recession were concentrated in ZIP codes with extensive subprime lending, while Atif Mian & Amir Sufi (2011) show that borrowing against appreciated home equity by existing homeowners was responsible for a significant fraction of the rise in U.S. household leverage from 2002 to 2006 (and the subsequent surge in defaults). Using instrumental variables estimation, they find that homeowners extracted 25 cents for every dollar increase in home equity, and that home equity-based borrowing added \$1.25 trillion in household debt from 2002 to 2008, potentially leading to a severe debt overhang depressing consumption in the following period.

However, more detailed sorting out of the determinants of the consumption bust that happened in the Great Recession is still in its infancy. For example, Ivaylo Petev, Luigi Pistaferri & Itay Saporta-Eksten (2001) use micro data from the Consumer Expenditure Survey and find that a decrease in consumption inequality in the Great Recession is largely explained by wealth shocks which hit the wealthy more than the poor. Karen Dynan (2012) uses micro data and shows that highly leveraged homeowners had larger declines in spending between 2007 and 2009 than other homeowners. Uncertainty also increased in the Great Recession, and Sule Alan, Thomas Crossley & Hamish Low (2012) demonstrate that a suitably calibrated model with credit constraints is able to explain the rise in the aggregate savings ratio in the UK during the Great

Recession if, in particular, young consumers faced a significant increase in uncertainty. The recession was also associated with depressed expectations for future income, and Mariacristina De Nardi, Eric French & David Benson (2012) show that a model with wealth shocks and shocks to consumers’ income expectations can explain the observed drop in consumption during the Great Recession.

The implications of a consumption bust can be large as found in a burgeoning theoretical literature that we will not review in detail; some examples are Gauti B. Eggertsson & Paul Krugman (2012), who demonstrate how debt overhang, affecting a large group of credit-constrained agents, can lead to stagnation resembling that observed in the western world following the 2007-2008 subprime crash, and Michael Kumhof, Romain Ranciere & Pablo Winant (2015), who model the interaction between household debt and income inequality, with “excess debt” triggering severe recessions.

3 Theoretical Background

We frame the discussion around a consumption model with housing. This model descends from the Permanent Income Hypothesis (PIH) of Hall (1978) and the buffer-stock model of Angus Deaton (1991), Christopher D. Carroll (1992), and Christopher D. Carroll (1997). Pierre-Olivier Gourinchas & Johnathan Parker (2002) find that U.S. consumers typically behave according to the buffer-stock model until about the age of 40, when consumption behavior changes to be more in accordance with the PIH due to accumulated life-cycle savings. However, in order to fully fit the data, important extensions are necessary, in particular, the existence of large illiquid assets—that is, housing—which generate large consumption commitments in the sense of Ray Chetty & Adam Szeidl (2007).

Consider the buffer-stock model with nondurables, owner-occupied housing, and down-payment requirements studied by María José Luengo-Prado (2006). Consumer j derives utility from the consumption of a nondurable good C and the services provided by housing H and maximizes expected util-

ity with respect to C and H :

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t U(C_{jt}, H_{jt}) \right\}, \text{ s.t. } S_{jt} = R_t S_{j,t-1} + Y_{jt} - C_{jt} - q_t \Delta H_{jt} - \chi(H_{jt}, H_{j,t-1}),$$

where the utility function typically is a CES index, S is financial wealth, q is the relative price of housing, R is an interest rate factor, and Y is income. There is a significant cost of moving, captured by the function $\chi()$, so that no consumer would adjust his or her housing stock marginally; i.e., consumers adjust their housing consumption only when their desired amount of housing (if there were no adjustment costs) significantly deviates from their current amount of housing.

The consumer further faces a collateral constraint, which limits borrowing to a fraction of the value of the housing stock. House-price appreciation is fully liquid for consumers for whom the collateral constraint is not binding; however, when house prices fall, many consumers will be unable to borrow because the debt limit binds. Consumers who suffer a temporary income shock may therefore end up cutting back disproportionately on non-durable consumption because it is not optimal to free up housing capital. This may make even wealthy individuals behave like they are constrained as in the models of “wealthy hand-to-mouth” consumers (Greg Kaplan, Giovanni Violante & Justin Weidner 2014) and “consumption commitments” (Chetty & Szeidl 2007). The debt limit itself is a function of personal income and credit scores, although a model with both these features seems not to have been studied quantitatively at this time. During the 2000s, the tightness of the constraint gyrated strongly, at least for subprime borrowers.

In simulations of the buffer-stock model, and of the just described housing model, log-income is typically assumed to be the sum of a random walk “permanent income” component and an i.i.d. transitory shock. If there is an above-average permanent income shock, consumers would like to increase consumption of both housing and nondurables, but the increase in consumption may be postponed while funds for the required down payment are accumulated. Costly foreclosure and geographical mobility can be added to the model

as in Yuliya Demyanyk, Dmytro Hryshko, María José Luengo-Prado & Bent E. Sørensen (2013).

3.1 Predicted consumption patterns

For easier reference in the empirical section, we provide a numbered list of the model’s “Consumption Predictions.”

1. Current and expected income growth cause current consumption growth in the PIH-model and in all forward-looking models that have followed. Consumption reacts one-to-one with permanent income shocks in the PIH-model of Robert E. Hall (1978) but less than that in the buffer-stock model of Christopher D. Carroll (2009), where the MPC is around 0.8 for standard parameterizations. Homeownership leads to lower MPCs even for permanent shocks, as demonstrated by Luengo-Prado (2006).
2. Higher uncertainty predicts lower current consumption in the buffer-stock model (Carroll 1992) and higher MPCs (also in aggregate data, María José Luengo-Prado & Bent E. Sørensen 2008). In the model, higher uncertainty can result from higher income variance, higher variance in house prices, or less risk sharing (which may or may not be reflected in measured income).
3. We do not have very clear predictions regarding concavity, given that the data are aggregated to the county level. Consumption is concave in transitory income shocks, with the strongest curvature around the point where the amount of liquid assets is equal to the desired buffer stock. (Deaton (1991); see also Alexander Michaelides (2003) for a comparison of Deaton’s and Carroll’s models.) The model does not have differential unconditional predictions regarding consumption growth rates for individuals with high or low permanent income. However, income fluctuates from year to year, and a fraction of high-income and low-income consumers in a given year will have transitorily high or low income—examples would be someone who wins the lottery or someone who has a

transitory illness; see the discussion in Milton Friedman (1957). For such people, income is dominated by a transitory component, and they will have lower MPCs than average consumers by the logic of the PIH; however, according to the buffer-stock model, consumers with transitory low income may have higher MPCs due to the concavity of the consumption function. Tullio Jappelli & Luigi Pistaferri (2014) find much lower MPCs for more affluent households using an Italian survey which directly asked about the response to transitory shocks. We look at income growth of high- and low-income percentiles even if the predictions of the models along this dimension are not clear for data such as ours.⁶ Mian, Rao & Sufi (2013) also consider concavity of the consumption function.

4. Tighter credit constraints will depress consumption growth because the desired buffer stock increases when the credit limit tightens. Sydney Ludvigson (1999) shows theoretically that a predictable tightening of credit limits leads to a decrease in consumption, while Thomas F. Crossley & Hamish W. Low (2014) empirically disentangle the direct effect (being currently credit constrained) from the indirect effect (accumulating a larger buffer stock of saving because credit will not be available if needed). They find a quarter of recent job losers, in a mid-1990s Canadian data set, was unable to borrow and therefore unable to smooth consumption. We expect the numbers to be similar in the United States although the severity of the constraints likely varied over our sample, where credit was eased in the subprime boom and then tightened in the Great Recession.
5. House prices are typically close to random walks (Wenli Li & Rui Yao 2007), as are stock prices (Eugene F Fama 1970). This implies that a housing wealth shock is equivalent to a transitory income gain.⁷ If

⁶Consumption functions do not need to be concave in wealth as shown by Luengo-Prado (2006). In the presence of housing and high credit limits, nondurable consumption functions are convex for low levels of wealth and concave for high levels of wealth.

⁷A permanent house-price shock is a one-time wealth shock equivalent to a temporary income shock. Only if the *growth* of house prices is integrated (rendering the house-price

homeowners have little wealth and the collateral constraint is binding, the house-price gain will be illiquid unless it is large enough.

6. An implication of the budget constraint is that net debtors will benefit from falling interest rates, while net savers will benefit from higher interest rates. In other words, debt will, everything else equal, predict increasing consumption in an environment of falling interest rates and *vice versa* in an environment of increasing interest rates. Benjamin J. Keys, Tomasz Piskorski, Amit Seru & Vincent Yao (2014) document, using micro data, a direct effect of mortgage interest rate resets on consumption.
7. In the PIH model, high debt is a reflection of expected high future income (John Y Campbell 1987); however, if these income gains do not materialize, as was the case for many in the Great Recession, debt will predict increased saving. Further, higher debt will predict lower consumption in the buffer stock model if repayments are higher than expected (lowering cash on hand) maybe because expected cash-out-financing becomes unavailable, as discussed in connection with credit. Kaplan, Violante & Weidner (2014) discuss how “wealthy hand-to-mouth” consumers may have significant, but illiquid, wealth and therefore high MPCs.
8. Expectations correlate with consumption. The less obvious issue is whether consumer expectations have predictive power that is not captured by other variables. Sydney C. Ludvigson (2004) finds that consumer confidence (which we interpret as a synonym for expectations regarding future real income) provides modest predictive power conditional on other observable variables. Christopher D Carroll, Jeffrey C Fuhrer & David W Wilcox (1994) find a similar result—they further find some evidence that consumer confidence may determine future income (via a multiplier effect). Robert B. Barsky & Eric R. Sims (2012) split expectations into a “news component” and an “animal spirits” compo-

itself twice integrated) will house-price shocks correspond to permanent income shocks.

nent and find the effect on future activity to be mainly related to the news component.

9. Foreclosure predicts falling consumption because it involves lack of access to credit, although this often involves slow erosion of credit ahead of the event and possibly a wealth shock; see Yuliya Demyanyk (2014) and Demyanyk et al. (2013).

Luengo-Prado & Sørensen (2008) show that the housing model can fit U.S. state-level MPCs well if significant risk sharing, as in Orazio P. Attanasio & Nicola Pavoni (2011), is added to the model. The standard one-good risk-sharing model (see Barbara J. Mace 1991, John H. Cochrane 1991) predicts that all consumers have perfectly coordinated consumption—a model which was rejected by, for instance, Cochrane (1991), Orazio Attanasio & Steven J. Davis (1996), and Dmytro Hryshko, María José Luengo-Prado & Bent E. Sørensen (2010) using micro data, and Pierfederico Asdrubali, Bent E. Sørensen & Oved Yosha (1996) and Yuliya Demyanyk, Charlotte Ostergaard & Bent E. Sørensen (2007) using regional data. Under imperfect risk sharing and correlated income shocks, consumption will be partly coordinated, and consumption patterns predicted at the individual level will survive aggregation as explicitly demonstrated by Sydney Ludvigson & Alexander Michaelides (2001) and Luengo-Prado & Sørensen (2008).

4 Data

We use multiple data sets. For growth variables, we calculate the growth rate over three years for each of the four periods: 2001–2003, 2004–2006, 2007–2009, and 2010–2012. For stock variables, we use the value in the year before the three-year period considered, with the exception of foreclosures (exact definition below). Most of our data are measured at the county-level, and we include all counties with a population over 5,000.

Consumption Growth. We use total retail sales at the county level from Moody’s to proxy for consumption. Total retail sales are the total of thirteen industries including both durables and nondurables: (1) motor vehicle and parts dealers, (2) furniture and home furnishings stores, (3) electronics and appliance stores, (4) building materials, garden equipment, and supplies dealers, (5) food and beverage stores, (6) health and personal care stores, (7) gasoline stations, (8) clothing and clothing accessories stores, (9) sporting goods, hobby, book, and music stores, (10) general merchandise stores, (11) miscellaneous store retailers, (12) non-store retailers, and (13) food services and drinking places. The data come from the Monthly Advance Retail Trade Survey (MARTS), which includes merchandise sold (for cash or credit at retail or wholesale) by establishments primarily engaged in retail trade. Services that are incidental to the sale of merchandise, and excise taxes that are paid by the manufacturer or wholesaler and passed along to the retailer, are also included. The monthly retail trade estimates are developed from samples representing all sizes of firms and kinds of business in retail trade, and the survey is composed of a sample selected from retail employers who made FICA payments.⁸

Income Growth. The county level data are from the Internal Revenue Service (IRS). We use real per capita adjusted gross income to construct three-year income growth rates. We are not able to estimate transitory versus temporary components of income, but the longer horizon filters out very short-lived shocks.⁹

Income Growth by Selected Percentiles. We calculate three-year income growth rates of the 90th and 25th income percentiles for each state and

⁸Retail sales include used cars, which are not typically included in units of cars sold, boats, motorcycles, recreational vehicles, parts, and repairs. Both retail and unit auto sales include fleet vehicle sales.

⁹Christopher D. Carroll & Andrew A. Samwick (1997) extract the permanent component of income, but we prefer not to impose stationarity of the statistical distributions of income across our very different sample periods. “Stationarity” in this context refers to the density of distributions, not to the time-series concept of stationarity—income is clearly not stationary in the time-series sense.

year from the Current Population Survey (CPS). This variable is available only at the state level. We do not consider lower percentiles than the 25th because these may be dominated by students and other individuals temporarily out of the labor force. (Preliminary estimations did not find significant results for those lower percentiles.)

Income Inequality. We calculate the share of income going to top 10 percent (the share of income of individuals earning more than the top 10th percentile in total income) using data from the CPS. This variable is available only at the state level.

Change in Unemployment Rate. We use data from the Bureau of Labor Statistics to construct the change in the county unemployment rate over the three years of each of the subperiods in our analysis.

Growth of Housing Wealth. We estimate real per capita housing wealth for counties in each year of our sample following the approach of Mian & Sufi (2011).¹⁰ Population and homeownership rates are county-level aggregates taken from the 2000 Census. House prices, measured by the house-price index (HPI), are available only for a subset of 1,245 unique counties. Whenever county-level data are not available, we substitute the missing observation for the county-level HPI with the corresponding state-level HPI.¹¹ Using this variable, we construct three-year growth rates.

Debt Overhang. To capture potential effects of debt on consumption growth, we use total debt at the beginning of the 3-year period and label it “debt over-

¹⁰For each county c and year t we calculate: $Housing\ Wealth_{c,t} = No.\ owner-occupied\ units_{c,t} \times Median\ home\ value_{c,t}$, where $No.\ owner-occupied\ units_{c,t} = No.\ owner-occupied\ units_{c,2000} \times \Delta Population \times \Delta Homeownership\ rate$. $Median\ home\ value_{c,t} = Median\ home\ value_{c,2000} \times (1 + \Delta House\ Prices_{c,(t-2000)})$, and $\Delta Population = (Population_t - Population_{2000})/Population_{2000}$. In this formula, the symbol “ Δ ” refers to the change since year 2000. $\Delta Homeownership\ rate$ is defined similarly to $\Delta Population$.

¹¹We verified that our results are virtually the same if we run regressions on the set of counties with non-missing county-level information on house prices.

hang.” We use individual-level data from Equifax available to us from the Consumer Credit Panel maintained by the Federal Reserve Bank of New York (“Equifax” for brevity hereafter) and aggregated over all individuals in each county to measure total debt. We use real per capita debt in the regressions but, in the descriptive section, we display 3-year growth rates. We use the lagged stock of debt in the regressions, rather than the change in debt, because the change in debt is a direct function of the change in consumption when income is controlled for, even though there is no exact one-to-one correspondence in our data.

Share of Subprime Borrowers. Individuals with credit scores below 661 are considered risky and usually referred to as “subprime borrowers.”¹² Such borrowers experienced a significant easing of access to credit during the subprime boom, with a reversal when the Great Recession broke. An easing of credit is likely to boost consumption of consumers with low credit ratings, in particular, and we interpret a significant coefficient on the subprime ratio as capturing a change in credit conditions, similarly to Mian & Sufi (2009). We use Equifax data to calculate the fraction of individuals in a county/year whose credit scores (Equifax Risk Scores) were below 661.

Fraction of Borrowers in Foreclosure. We calculate this fraction as the number of consumers who had at least one foreclosure in the past 24 months relative to the number of all consumers in Equifax by county and year. The choice of the past 24-month period is dictated by data availability. Because of the backward looking nature of the raw data, this variable is measured at the end of the period (i.e., for each period $t - 3$ to t foreclosure is measured as of period t).

Consumer Expectations. We use data on consumer expectations from the Conference Board, available for nine Census Divisions, which we match with the counties in our sample. Our index of expectations is the average of three

¹²See, for example, <http://investor.equifax.com/releasedetail.cfm?ReleaseID=881777>

indexes which measure consumers' perceptions about business conditions, employment conditions, and total family income six months hence. The data are monthly and we average to the annual frequency before calculating three-year growth rates.

Share of Cash-Out Refinancings. We calculate this share as the number of cash-out refinance originations scaled by the number of outstanding mortgages, and use its three-year growth rate in our regressions. The data are from Black Knight Financial Services, Inc.

5 Descriptive Statistics

The empirical analysis uses county-level data. There is large variation in consumption growth, income growth, etc. across counties, even within states, but in order to provide uncluttered illustrations, most figures depict state-level patterns.

Figure 1 shows the growth rates of (real per capita) aggregate U.S. total, nondurable, durable, and services expenditures together with the growth rates of county-level retail sales aggregated to the U.S. economy. Nondurable consumption grew modestly during the dot-com recession, at about 1 percent, accelerated to over 8 percent during the subprime boom, fell in the Great Recession, and grew by about 7 percent in the tepid recovery. This latter number is not as impressive as it sounds; in fact, it is atypical of U.S. business cycles for consumption to not have fully recovered this long after a recession; see Petev, Pistaferri & Saporta-Eksten (2001). Expenditure on durables fell particularly strongly in the Great Recession, by an astonishing 21 percent. It increased in the tepid-recovery but, as for most components, the increase was tepid in the sense that it did not make up for lost ground in the Great Recession. The strong collapse in durables is consistent with the model of Martin Browning & Thomas F. Crossley (2009). Services was one of the fastest growing components in the dot-com recession and the subprime boom, but the consumption of services is virtually unchanged since then. Total consumption

was less volatile than its components.

Goods is the combination of nondurables and durables. Overall, retail sales match goods consumption quite well. For example, the drop in retail sales in the Great Recession was about 13 percent while goods consumption dropped about 10 percent. The difference between retail sales and goods consumption is smaller in the other periods. Our regressions are cross-sectional and focus on the relative importance of consumption determinants across counties, but it is reassuring that the growth rates are similar in the aggregate.

Figure 2 provides evidence of cross-county variation in consumption growth rates in a box-and-whisker plot, where the top and bottom of the boxes are the 75th and 25th percentiles, respectively. The data for this plot (and our regressions) is winsorized at 2 and 98 percent.¹³ The interquartile ranges span about 10 percentage points in each period, and some counties have consumption growth rates that are far different from those of other counties as evidenced by some county-observations falling outside the “whiskers.”¹⁴ The counties with atypical growth rates are mostly counties with relatively high growth rates in the recessions and relatively low growth rates in the subprime boom and the tepid recovery; even during the subprime boom when aggregate consumption grew at a fast pace of 6.1 percent per year, some counties had negative growth rates of over 20 percent. Natural disasters, such as hurricane Katrina, which hit the Gulf Coast and, in particular, New Orleans, in 2005, generated large negative outliers which will have undue influence in the absence of winsorizing.

Our data provide further details not readable from the figure: in the Great Recession, 2,618 out of 2,768 counties had consumption growth less than 5 percent, while 1,050 counties (about a half of all counties in the United States) experienced a decline larger than 15 percent. The tepid recovery was not uniformly distributed either: twenty states had weak consumption growth (positive growth rates smaller than 8 percent), while consumption grew quite rapidly at rates above 8 percent in the remaining states.

¹³A similar plot of the non-winsorized data is provided in the Appendix.

¹⁴The length of the whiskers is 1.5 times the interquartile range.

Figure 3 provides a map of the U.S. states indicating the geographical distribution of consumption growth rates (retail sales). During the dot-com recession, 25 states had negative consumption growth while, during the subprime boom, only Michigan, likely due to contraction in the car industry, had negative three-year consumption growth. During the Great Recession, all states had negative consumption growth, but the decline was not uniform. One state had consumption growth between 0 and -5 percent, 4 states between -5 and -10 percent, 26 states between -10 and -15 percent, and a staggering 20 states had consumption falling by more than 15 percent.

Figure 4 shows the distribution of changes in the unemployment rate and the growth rate of state-level income, debt overhang, and consumer expectations. Comparing Figure 3 and Panel A of Figure 4, we see that while some states had negative growth in both consumption and income during the dot-com recession, five states had rising income but declining consumption. All states experienced a sharp fall in consumption in the Great Recession, but income did not show the same pattern. During this period, 35 states had positive income growth, and five states among them had real per capita income growth of more than 8.5 percent. In the tepid recovery, income growth was high for a large fraction of states: 23 states had income growth larger than 8.5 percent.

Consumer “overborrowing” during the subprime boom is sometimes blamed for the severity of the recession that followed and for the slow speed of the tepid recovery (e.g, Mian & Sufi 2009, Dynan 2012). In Panel B of Figure 4, all but one state had debt growing by more than 10 percent during the dot-com recession while, during the subprime boom, 29 states had debt growing by more than 10 percent. By the time consumption plummeted in all states and counties in the Great Recession, only 12 states deleveraged to an extent that their debt was shrinking on average; however, in the tepid recovery debt was shrinking in all states except North Dakota, which was booming due to rapidly expanding oil production. Further, although not visible in the figure, 36 states deleveraged by more than 10 percent, of which 16 states deleveraged by more than 15 percent.

The trends in unemployment are even more pronounced; see Figure 4, Panel C. In the dot-com recession, unemployment increased more than 1.5 percentage points in 32 states, mainly outside the Southeast and the Rocky Mountains. In the subprime boom, almost all states outside of the Midwest increased employment while, in the Great Recession, every single state had increasing unemployment. During the tepid recovery, the change in unemployment was quite scattered across states.

Changes in consumer expectations (see Figure 4, Panel D) were small during the dot-com recession, indicating that the recession felt relatively mild to consumers, but the picture was drastically different in the Great Recession, when consumer expectations collapsed by more than 26 percent in all states except those in the New England Census Division. Consumer expectations improved across the board during the subprime boom and the tepid recovery.

In Figure 5, Panels A and B, we plot the growth rate of state-level real per capita income of the 90th and the 25th percentiles of the overall income distribution. The figure shows no uniform trends in income growth of the poor versus the rich. Income of the lower percentile held up better than income at the top in the Great Recession, with many states displaying positive growth of the 25th percentile, which may be the result of federal stimulus helping the less advantaged, while households in the higher income brackets saw steep declines in capital income during this period.

Panel C of Figure 5 displays variation in housing wealth across states and periods. For housing wealth the difference between the two recessions in our sample was dramatic: in the dot-com recession, states either had rapidly growing or fairly constant housing wealth while, in the Great Recession, no state had significant growth in housing wealth and 14 states had housing wealth declining by more than 15 percent. During the tepid recovery, 35 states had housing wealth declining by more than 15 percent.

States with large fractions of subprime borrowers are mostly concentrated in the South. These fractions are more stable over time than any other measure we used in our analysis; see Figure 5, Panel D.

We use the observed variation in all county- or state-level variables listed

and an array of other variables not included in the figures to assess the *contribution* of each factor in explaining the growth of county-level consumption during the four subperiods considered.

6 Specification and Results

We believe that the volatility of the 2000-2012 period precludes the use of panel-data estimations with all years pooled, and we therefore split the sample into four periods. One possibility is to use regressions with lags in order to pinpoint the exact pattern of adjustment to shocks, but we are interested in testing the impact on consumption of the variables suggested in the literature in a nested framework. This involves a large number of regressors, and dynamic regression would then entail a very large number of estimated coefficients, leading to a confusing and unstable picture. Instead we follow, for example, Emi Nakamura & Jon Steinsson (2014) and Petev, Pistaferri & Saporta-Eksten (2001), and use long time intervals with three-year growth rates. Such regressions, where the constant captures the aggregate growth rate of consumption, can also be seen as measuring which variables generate deviations from perfect risk sharing.¹⁵

We regress consumption growth on an array of macroeconomic variables such as income, unemployment, debt, income inequality, consumer confidence, housing wealth, access to credit, and foreclosures. For robustness, we run regressions on a variety of other variables to make sure we do not omit any important determinants of consumption growth. The variables that do not explain consumption growth above and beyond those listed above are not included in the paper for brevity.

We do not have instruments which would allow us to give clear causal results of, say, a change in the unemployment rate separate from changing income expectations. For example, Thomas Philippon & Virgiliu Midrigan

¹⁵Cochrane (1991) and Hryshko, Luengo-Prado & Sørensen (2010) estimate the impact on risk sharing of unemployment and house-price appreciation, respectively, using long time intervals.

(2011) construct a model where a tightening of credit constraints causes declining consumption and increased unemployment, so that unemployment is a result of the credit tightening and in a structural sense “everything” is caused by credit tightening. Similarly, house-price growth may be a function of income growth, which itself may be a function of aggregate demand shocks or of productivity supply shocks, and we are not able to sort this out. We will however talk about the “effect” of each variable without torturous attempts to signal this caveat in further writing—our descriptive approach does not uncover deep structural forces but rather provides stylized facts.

We estimate the following cross-sectional regressions over U.S. counties:

$$\Delta^3 \log(C_{c,t}) = \alpha + \beta_1 \widetilde{\Delta^3 \log(X_{c,t})} + \beta_2 \widetilde{Y_{c,t-3}} + \beta_3 \widetilde{\Delta^3 UR_{c,t}} + \beta_4 \widetilde{Foreclosure_{c,t}} + \epsilon_{ct}, \quad (1)$$

where $\Delta^3 \log(C_{c,t}) = \log(C_{c,t}) - \log(C_{c,t-3})$ is the three-year growth rate of county-level consumption proxied by real per capita total retail sales, $\Delta^3 \log(X_{c,t})$ is the three-year growth rate of county-level variables (or state- or census region-level variables for which county-level data are not available), $\Delta^3 UR_{c,t} = UR_{c,t} - UR_{c,t-3}$ is the change in the unemployment rate over the subperiod, and $Foreclosure_{c,t}$ is the share of consumers who had at least one foreclosure in the past 24 months relative to the number of all consumers, measured at the end of the period. $Y_{c,t-3}$ are county-level lagged variables—we prefer to include the lagged value for most of our stock values. In particular, debt is a one-to-one function of consumption, everything else equal, and although this identity does not hold in our data where consumption is approximated by retail sales, consumption shocks would be reflected in debt. Other stock variables, such as the income share of the top 10 percent, are slowly changing, and the interpretation is cleaner when using the predetermined value.

We demean all independent variables in order for the constant to capture average consumption growth over each three-year interval in the following way: $\widetilde{\Delta \log(X_{c,t})} = \Delta \log(X_{c,t}) - \frac{1}{N} \sum_{c=1}^N \Delta \log(X_{c,t})$, where c indexes counties and N is the total number of counties in our sample. Lagged and concurrent variables are demeaned in a similar way. Our data have significant outliers

(see Figure A.1 in the Appendix) and we therefore winsorize all variables at 2 and 98 percent to make sure our results are not driven by outliers. Most results are quite robust to winsorizing but, for example, the estimated effect of income growth varies more across subperiods with non-winsorized data. Standard errors were estimated robustly. Our main results are presented in Table 2. We organize the discussion of the results by regressor.

Constant. The constant measures average county-level consumption growth because the regressors are demeaned. Consumption declined weakly (−1 percent) over the dot-com recession, recovered by 7 percent over the subprime boom, and fell steeply (−12 percent) during the Great Recession, highlighting how much more severe this recession was than the dot-com recession. During the tepid recovery, consumption grew by 10 percent, almost recovering the ground lost in the Great Recession but, as pointed out by Petev, Pistaferri & Saporta-Eksten (2001), it is atypical for consumption to be depressed that long after the onset of a recession. These patterns are driven by aggregate effects, for example, the drop in the Great Recession is consistent with a U.S.-wide increase in uncertainty, a drop in income expectations, and a loss of wealth; however, without more degrees of freedom, we cannot test this and we move on to the estimated determinants of county differences. The effects of most economic variables differ between periods.

Income. The elasticity of per capita consumption with respect to per capita income is quite robust across periods at about 10 percent, with high statistical significance in every period. Income is of course correlated with many variables, and the coefficient is about 20 percent in untabulated regressions of consumption on income alone. Compared to Consumption Prediction # 1, this is quite a low value, which may indicate that county-level income shocks are considered transitory by consumers. Further, the IRS income data are likely to be a noisy measure of labor income, and this may partly explain the very low elasticities found here. Cross-county shopping, and substitution of state-level variables for some county-level variables, might also add to a downward bias in the elasticities.¹⁶

¹⁶Luengo-Prado & Sørensen (2008) find MPCs around 0.33 for nondurable state-level

Unemployment. For a consumer, job loss is typically associated with a large negative income shock. However, our regressions control for income, and our preferred interpretation of unemployment, in the context of the model, is that high unemployment in a county is associated with high income uncertainty, Consumption Prediction # 2. The effect of rising or falling unemployment is also estimated with high stability across periods and with even higher precision than was found for income. The economic interpretation of the coefficient of -0.01 is that a 1 percentage point increase in unemployment will decrease consumption by 1 percent. Clearly, changing unemployment, whether increasing or decreasing, was a strong predictor of consumption throughout.

Growth of Income Percentiles. Income percentiles are available only at the state level and average income is not highly correlated with the percentiles. The coefficients to the percentiles can therefore be interpreted as incremental effects over average income. In the dot-com recession, the distribution of income mattered, with higher consumption growth if income growth of the 90th and 25th percentiles was higher—with the latter estimate being more significant. Less wealthy individuals may have higher propensities to consume—Consumption Prediction # 3—as typical models of consumption, such as Carroll (1992), predict a concave pattern in wealth normalized by permanent income. Because our time series are too short to estimate permanent income, we cannot literally test this prediction, but it is still of interest to consider non-linearities in income. In the tepid recovery, the expected concave pattern is found with consumption reacting significantly more to growth of the 25th percentile of income and less to growth of the 90th percentile. Overall, these patterns are not very strong in our data and they will be challenging to model: concavity of the consumption function in wealth can likely be captured by standard consumption models with suitable distributions of shocks to wealth and permanent income, while the U-shaped pattern in the dot-com recession presents a challenge. A model with a sharp loosening of credit constraints, which allows the less wealthy to save for and purchase real estate,

retail sales during 1970–1998, which they were able to match using the model with housing described in Section 3 when adding substantial (not directly measured) risk sharing.

may be able to explain periods such as the subprime boom, but what we can say with some confidence is that a representative-agent assumption likely will lead to invalid results for the period between 2000 and 2012.

Inequality: Income share of the top 10 percent. High-income consumers typically have high wealth (which we have no precise direct measures of) and might be able to better withstand income shocks in the sense of keeping a higher level of nondurable consumption by adjusting asset holdings. We find that consumption indeed falls less in wealthy counties in both recession periods, while the point estimates are small and insignificant in the subprime boom and in the tepid recovery. More studies are needed to uncover whether this pattern is due to wealth or, e.g., to wealthy individuals' income being less volatile, although the results of Petev, Pistaferri & Saporta-Eksten (2001) indicate that the wealthy did not fare that well in the Great Recession. Having controlled for total income and income of the 90th percentile, we believe this variable captures lower uncertainty, Consumption Prediction # 2, for the well-to-do, who typically are not working in sectors, such as construction, where the probability of job loss is high in recessions.

Housing Wealth. Mian, Rao & Sufi (2013) find large effects of housing wealth on consumption in the Great Recession, and we confirm this result, finding strong positive effects on consumption growth in the subprime boom and strong negative effects in the Great Recession. There is a marginally significant positive effect in the dot-com recession, implying that house-price growth helped stabilize the economy during that period, while housing appreciation helped fuel the subprime boom and contributed to the subsequent Great Recession. In the tepid recovery, the rebound in house prices had little effect. We conjecture that this is a reflection of many homeowners being underwater—owing more on their mortgage than their house is worth—and/or tighter credit constraints such that the wealth gain from any house-price increase was illiquid and could not be borrowed against. According to Consumption Prediction # 5, the propensity to consume out of increasing housing wealth should be small, so the magnitude of the coefficients in the subprime boom and the Great Recession may indicate that consumers expected an ongoing house-price ap-

preciation.

Debt overhang. Debt overhang has a clear negative relationship with consumption growth except in the dot-com recession. (Dynan (2012) found negative effects of debt overhang in the Great Recession using micro data.) Our interpretation is that the interest rate declines during the dot-com recession made debt service less burdensome (Consumption Prediction # 6) which allowed indebted consumers to increase consumption, as in the textbook model of interest rate effects on net debtors—the change in (conventional, first lien) mortgage rates over the subperiods is displayed in Figure 6. During the subprime boom, interest rates increased, and the burden of carrying debt became heavier, depressing consumption. When the recession hit and credit conditions became much tighter, consumption was depressed due to credit constraints, even as interest rates continued to decline. We interpret this pattern as reflecting that a large number of consumers needed to pay down debt and limit consumption. Consumers with good credit and stable jobs were able to re-finance into lower interest rates and increase consumption, but implicit in Consumption Prediction # 7 is that life-cycle consumer debt reflects expected future savings (a function of income) so that a downward revision of expected future income makes debtors increase their current saving rate. Consumption Predictions # 6 and # 4, concerning interest rates and credit, are still relevant but it appears that the wealth gains of debtors from lower interest rates were dominated by the credit and income effects in the Great Recession. Consumer deleveraging continued into the tepid recovery contributing to the slow speed of recovery.

Subprime Credit. As a direct indicator of credit conditions, Consumption Prediction # 4, we include the share of subprime borrowers; i.e., the fraction of individuals whose Equifax credit scores is lower than 661. We interpret this measure as capturing changes in the amount of credit available to marginal borrowers with impaired credit. The subprime boom period stands out with the fraction of subprime borrowers being an important predictor of consumption during those years. This finding strongly agrees with the results of Mian & Sufi (2009)—they mainly consider home equity lending in isolation, but we

show that the easing of credit has strong significance even after all our other variables have been included.

Cash-Out Refinancings. Another indicator of credit conditions is the share of cash-out refinancings. Borrowers with good credit were able to refinance into lower rates in the tepid recovery, but this was not possible for unemployed individuals or individuals with bad credit, underwater mortgages, or other impediments. Maybe not surprisingly, this variable is highly significant only in the tepid recovery. (We alternatively considered the number of individuals with very high credit scores, but this was not significant—likely because a good credit score was a necessary but not a sufficient condition for being able to refinance in the tepid recovery.) The theoretical underpinnings for the importance of this variable are found in Consumption Predictions # 6 and # 4, regarding interest rates and credit.

Foreclosure. A foreclosure is costly and limits future access to credit severely, as summarized in Consumption Prediction # 9. In the run-up to foreclosure, many consumers may cut back consumption, hoping to avoid it and, indeed, the number of foreclosures in a county correlates negatively with consumption growth in all subperiods except the Great Recession. During the Great Recession, a large number of foreclosures was a result of collapsing house prices, and it seems that falling house prices during this period leave no further explanatory power for foreclosure.

Consumer Confidence. Expectations are at the core of rational expectations consumer models, Consumption Prediction # 8. Even if our measure of expected economic performance is available for only 9 Census Divisions, it is significant in every subperiod, except the tepid recovery, confirming that consumers act on their expectations and increase consumption more when economic conditions are expected to improve. It is virtually impossible to know if the effect is due to rational forecasting of expected future income or fear of some negative tail-event.

The variables in our statistical model are able to explain between 5 and 13 percent of the county variation in consumption growth, based on the adjusted R^2 . Retail sales are an imperfect measure of consumption and, due

to cross-county shopping, retail sales in a county may reflect consumption in neighboring counties. Such noise will tend to depress our explanatory power and will make it hard to find significant effects. These imperfections, however, simply add to the error term and, because they are unlikely to be systematically correlated with our regressors, do not lead to bias. The total R^2 was particularly low in the dot-com recession and the tepid recovery.

To assess the economic significance of the results, we plot in Figure 7 (i) the average values of each variable used in our estimation, (ii) the estimated coefficient multiplied by one standard deviation of the variable (calculated cross-sectionally for each relevant sample period) indicating economic significance, and (iii) partial R^2 s for each variable showing the share of variance explained by each variable.

6.1 Economic Significance

We illustrate the economic significance of the regressors in Figure 7. The left panel shows the average value of the regressors across all counties for each period. While this average value technically does not help pin down our cross-sectional estimates, it helps paint a picture of the variation over our sample periods, and together with the estimated parameters is indicative of aggregate effects. Over the periods considered, per capita income growth varies from a low of -6.6 percent during the Great Recession to a high of 12.3 percent during the tepid recovery. The middle panel of Figure 7 is directly informative of impact of each variable in explaining differences in consumption growth across counties as it shows the estimated coefficient times the cross-sectional standard deviation of the relevant regressor in each period. The effect of income growth is fairly stable over time: a county with a one-standard-deviation higher income growth than average is predicted to have 0.5 – 0.8 percent higher consumption growth. A one-standard-deviation change in unemployment explained more than a percentage point of the change in consumption in each period. Growth of the 90th and 25th percentiles is less economically significant than average growth, but the picture adds to our understanding. In particular, income of

the least affluent fell steeply in the tepid recovery, which may help explain why consumption only slowly rebounded.¹⁷ Our inequality measure is fairly constant over time and has some economic significance in the recessions but not in the boom or the tepid recovery.

Not surprisingly, changes in housing wealth contributed significantly to consumption patterns in the subprime boom and the Great Recession. Debt overhang, while quite stable, predicted large negative changes in consumption in all periods except the dot-com recession. Share of subprime borrowers added significantly to consumption growth in the subprime boom: while the average fraction of subprime borrowers is fairly stable over time, it is important for explaining variation in consumption across counties during the subprime boom, but not in other subperiods. A county with a one-standard-deviation higher proportion of such borrowers had a 1 percent higher consumption growth rate in the subprime boom, while the effect is substantially lower and varies from 0.1 to 0.3 percent during other periods. The total amount of debt at the beginning of the period, while not significantly different over these four periods, is associated with consumption growth during the dot-com recession, but in the other periods debt is associated with shrinking consumption. Debt contributes greatly to explaining the variation in consumption growth across counties, although its importance has been declining. A one-standard-deviation higher debt at the beginning of the period increases consumption by 1.3 percent in 2004–2006, while it lowers consumption by 0.9 percent by 2009–2012.

Foreclosures had a negative impact on consumption outside of the Great Recession—possibly foreclosures were important in the Great Recession but being across the board, our regressions cannot capture this; likely, micro data are needed to fully sort this out. Changes in consumer expectations, which vary greatly over time, play a strong role during the subprime boom where expectations were positive: consumption is about 1.3 percent higher in regions with one-standard-deviation more optimistic expectations than the average.

¹⁷The left-hand column depicts averages across countries, while our results are identified from the deviation from these averages, so our interpretations of the numbers in the middle panel of Figure 7 are only suggestive.

The estimated impact was smaller during other periods, even if expectations were quite negative in both recessions. Finally, counties with a large share of cash-out refinancings—which we take to be an indicator of excellent access to credit—did particularly well in the tepid recovery. One standard deviation increase in cash-out refinancings contributed to 0.8 percent higher consumption growth during this period.

6.2 Marginal Explanatory Power: Partial R^2 Analysis

In the third panel of Figure 7, we plot the partial R^2 for each variable in order to assess how much each variable is contributing to the model fit. The partial R^2 s (multiplied by 100) and the contribution to the R^2 (the partial R^2 as a share of the total R^2) are displayed in Table 3.¹⁸ In the dot-com recession and the tepid recovery, debt overhang and unemployment have the most predictive power—with partial R^2 s contributing about 20 percent each of the total R^2 . It appears that monetary policy was particularly effective in the dot-com recession—with falling interest rates and easy credit, debtors could refinance into cheaper loans and benefit from the increase in debtors’ effective wealth.

Overall, debt overhang and unemployment growth are the most important determinants of consumption growth throughout. During the subprime boom, the fraction of subprime borrowers was also important, while inequality was important in the dot-com recession. In the tepid recovery, income growth, foreclosures, and cash-out refinancings were all important.

7 Conclusion

We explain the variation in consumption growth across U.S. counties during the first twelve years of the millennium. Using a rich data set, we document

¹⁸These partial R^2 s do not sum to the total R^2 unless the regressors are orthogonal because each of them only measures the incremental explanatory effect of the relevant variable so, for example, if all the regressors were highly correlated, each partial R^2 would be very small.

the explanatory power of numerous economic variables during each of four subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). We find that income, house prices, unemployment, consumer confidence, and cash-out refinancings as well as the (lagged) level of debt, inequality, the fraction of borrowers that are subprime, and the fraction of mortgages in foreclosure help explain consumption growth during our sample period, albeit differently in different subperiods.

Unemployment variation and debt overhang have the largest explanatory power for consumption growth throughout our sample. Optimistic consumer expectations were important in the subprime boom, while housing wealth was important in the subprime boom and in the Great Recession. The fraction of subprime borrowers significantly predicted consumption growth in the subprime boom, and income growth, foreclosures, and cash-out refinancings were important in the tepid recovery. Overall, many of the patterns found are specific for certain subperiods, and any full modeling of the boom and bust of the 2000s needs to account for these complicated patterns.

We do not calibrate a structural model, but our results provide important facts for models to match. Several factors in our data likely are causal for consumption growth while others are not. For example, debt overhang is likely to be a predetermined causal determinant, while unemployment may capture income uncertainty, income expectations, or more complicated interactions or nonlinearities. We find that income growth is correlated with consumption growth in all four subperiods, and we believe this is a direct causal relation because feedback from consumption to income is likely to be minor: for example, the consumption component captured by retail sales will typically be mainly produced out-of-county (as opposed to construction and restaurant meals that are not included). Growth of housing wealth is also important in most periods, while counties with high inequality (high income share of the top 10 percent) performed better in the two recessions in our sample. Indicators of credit availability were significant in the subprime boom and the tepid recovery, although our regressions cannot rule out that these patterns

capture, for example, income expectations. Consumer expectations predicted consumption growth in all periods except during the tepid recovery.

Our study contributes to a large body of literature that either empirically uncovers or models determinants of consumption growth during the 2000s. Our main contribution lies in quantifying the relative impact of a variety of factors that affected the economy during this period. Further work might include the collection of even more detailed data in order to address issues of aggregation, and possibly allow for more rigorous identification of causality, and the formulation of models calibrated to the facts uncovered here.

References

- Alan, Sule, Thomas Crossley, and Hamish Low.** 2012. “Saving On a Rainy Day, Borrowing For a Rainy Day.” IFS Working Paper 1211.
- Asdrubali, Pierfederico, Bent E. Sørensen, and Oved Yosha.** 1996. “Channels of Interstate Risk Sharing: United States 1963–1990.” *Quarterly Journal of Economics*, 111: 1081–1110.
- Attanasio, Orazio, and Steven J. Davis.** 1996. “Relative Wage Movements and the Distribution of Consumption.” *Journal of Political Economy*, 104(6): 1227–1262.
- Attanasio, Orazio P., and Nicola Pavoni.** 2011. “Risk Sharing in Private Information Models With Asset Accumulation: Explaining the Excess Smoothness of Consumption.” *Econometrica*, 79(4): 1027–1068.
- Barsky, Robert B., and Eric R. Sims.** 2012. “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence.” *American Economic Review*, 102(4): 1343–77.
- Browning, Martin, and Thomas F. Crossley.** 2009. “Shocks, Stocks, and Socks: Smoothing Consumption Over a Temporary Income Loss.” *Journal of the European Economic Association*, 7(6): 1169–1192.
- Campbell, John Y.** 1987. “Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis.” *Econometrica*, 55(6): 1249–73.
- Carroll, Christopher D.** 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *Brookings Papers on Economic Activity*, 23(2): 61–156.
- Carroll, Christopher D.** 1997. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *The Quarterly Journal of Economics*, 112(1): 1–55.

- Carroll, Christopher D.** 2009. "Precautionary saving and the marginal propensity to consume out of permanent income." *Journal of Monetary Economics*, 56(6): 780–790.
- Carroll, Christopher D., and Andrew A. Samwick.** 1997. "The Nature of Precautionary Wealth." *Journal of Monetary Economics*, 40: 41–71.
- Carroll, Christopher D, Jeffrey C Fuhrer, and David W Wilcox.** 1994. "Does Consumer Sentiment Forecast Household Spending? If So, Why?" *American Economic Review*, 84(5): 1397–1408.
- Chetty, Ray, and Adam Szeidl.** 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics*, 122(2): 831–877.
- Cochrane, John H.** 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy*, 99(5): 957–976.
- Crossley, Thomas F., and Hamish W. Low.** 2014. "Job Loss, Credit Constraints, and Consumption Growth." *Review of Economics and Statistics*, 96(5): 876–884.
- Deaton, Angus.** 1991. "Saving and Liquidity Constraints." *Econometrica*, 59: 1121–1248.
- Demyanyk, Yuliya.** 2014. "The Impact of Missed Payments and Foreclosures on Credit Scores." FRB Cleveland Working Paper 14-23.
- Demyanyk, Yuliya, and Otto Van Hemert.** 2011. "Understanding the Subprime Mortgage Crisis." *Review of Financial Studies*, 24(6): 1848–1880.
- Demyanyk, Yuliya, Charlotte Ostergaard, and Bent E. Sørensen.** 2007. "U.S. Banking Deregulation, Small Businesses, and Interstate Insurance of Personal Income." *Journal of Finance*, 62(6): 2763–2801.
- Demyanyk, Yuliya, Dmytro Hryshko, María José Luengo-Prado, and Bent E. Sørensen.** 2013. "Moving to a Job: The Role of Home Equity, Debt, and Access to Credit." CEPR Discussion Paper 9474.

- De Nardi, Mariacristina, Eric French, and David Benson.** 2012. "Consumption and the Great Recession." *Economic Perspectives*, , (1Q): 1–16.
- Dynan, Karen.** 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity*, 44(1): 299–362.
- Eggertsson, Gauti B., and Paul Krugman.** 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *The Quarterly Journal of Economics*, 127(3): 1469–1513.
- Fama, Eugene F.** 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, 25(2): 383–417.
- Friedman, Milton.** 1957. *A Theory of the Consumption Function*. Princeton University Press.
- Gourinchas, Pierre-Olivier, and Johnathan Parker.** 2002. "Consumption over the Life Cycle." *Econometrica*, 70: 47–91.
- Hall, Robert E.** 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86: 971–987.
- Hryshko, Dmytro, María José Luengo-Prado, and Bent E. Sørensen.** 2010. "House Prices and Risk Sharing." *Journal of Monetary Economics*, 57(8): 975–987.
- Iacoviello, Matteo.** 2011. "Housing Wealth and Consumption." International Finance Discussion Paper 1027.
- Jappelli, Tullio, and Luigi Pistaferri.** 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics*, 6(4): 107–36.
- Kaplan, Greg, Giovanni Violante, and Justin Weidner.** 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*, , (Spring): 77–138.

- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vincent Yao.** 2014. "Mortgage Rates, Household Balance Sheets, and the Real Economy." NBER Working Paper 20561.
- Kumhof, Michael, Romain Ranciere, and Pablo Winant.** 2015. "Inequality, Leverage, and Crises." *American Economic Review*, 105(3): 1217–45.
- Li, Wenli, and Rui Yao.** 2007. "The Life-Cycle Effects of House Price Changes." *Journal of Money, Credit and Banking*, 39(6): 1375–1409.
- Ludvigson, Sydney.** 1999. "Consumption And Credit: A Model Of Time-Varying Liquidity Constraints." *The Review of Economics and Statistics*, 81(3): 434–447.
- Ludvigson, Sydney, and Alexander Michaelides.** 2001. "Can Buffer Stock Saving Explain the Consumption Excesses?" *American Economic Review*, 91: 631–647.
- Ludvigson, Sydney C.** 2004. "Consumer Confidence and Consumer Spending." *Journal of Economic Perspectives*, 18(2): 29–50.
- Luengo-Prado, María José.** 2006. "Durables, Nondurables, Down Payments and Consumption Excesses." *Journal of Monetary Economics*, 53(7): 1509–1539.
- Luengo-Prado, María José, and Bent E. Sørensen.** 2008. "What Can Explain Excess Smoothness and Sensitivity of State-Level Consumption?" *Review of Economics and Statistics*, 90(1): 65–80.
- Mace, Barbara J.** 1991. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy*, 99(5): 928–956.
- Mian, Atif, and Amir Sufi.** 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *The Quarterly Journal of Economics*, 124(4): 1449–1496.

- Mian, Atif, and Amir Sufi.** 2011. “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis.” *American Economic Review*, 101(5): 2132–56.
- Mian, Atif, Kamalesh Rao, and Amir Sufi.** 2013. “Household Balance Sheets, Consumption, and the Economic Slump.” *Quarterly Journal of Economics*, 128(4): 1987–1726.
- Michaelides, Alexander.** 2003. “A Reconciliation of Two Alternative Approaches Towards Buffer Stock Saving.” *Economics Letters*, 79(1): 137–143.
- Nakamura, Emi, and Jon Steinsson.** 2014. “Fiscal Stimulus in a Monetary Union: Evidence from US Regions.” *American Economic Review*, 104(3): 753–92.
- Petev, Ivaylo, Luigi Pistaferri, and Itay Saporta-Eksten.** 2001. “Consumption and the Great Recession: An Analysis of Trends, Perceptions and Distributional Effects.” In *Analyses of the Great Recession.* , ed. David Grusky, Bruce Western and Christopher Wimer, 161–195. New York:Russell Sage Foundation.
- Philippon, Thomas, and Virgiliu Midrigan.** 2011. “Household Leverage and the Recession.” NBER Working Paper 16965.

Figure 1: Growth of U.S. Retail Sales and Consumption Components

This figure compares three-year growth rates of real consumption components and aggregated total retail sales, calculated as $\Delta^3 \log(C_t) = 100 \times (\log(C_t) - \log(C_{t-3}))$ for each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). The growth rate of total personal consumption is labeled *Consumption*; its two sub-components are *Goods* and *Services*. *Goods* is the sum of *Durables* and *Nondurables*. *Durables* consist of personal expenditures on motor vehicles and parts, furnishings and durable household equipment, recreational goods and vehicles, and other durable goods. *Nondurables* are goods in the following categories: food and beverages purchased for off-premises consumption, clothing and shoes, gasoline, fuel oil and other energy goods, and other nondurable goods. We also plot the growth rates of *Services* which consist of household consumption expenditures, housing and utilities, health care, transportation, recreation, food services and accommodations, financial services and insurance, and other services. The data source is Moody's Analytics and the Bureau of Economic Analysis.




Figure 2: County Retail Sales Growth

This figure displays three-year growth rates of real per capita county-level consumption growth, proxied by total retail sales, calculated as $\Delta^3 \log(C) = 100 \times (\log(C_t) - \log(C_{t-3}))$ for each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). The data source is Moody's Analytics. The data are winsorized at the 2 and 98 percent level.




Figure 3: State Consumption Growth

This figure displays three-year growth rates of real per capita consumption proxied by the total county-level retail sales aggregated to the state level and calculated as $\Delta^3 \log(C_t) = 100 \times (\log(C_t) - \log(C_{t-3}))$ for each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). The data source is Moody's Analytics.




Figure 4: Change in State Unemployment Rate and Growth Rates of State Income, Debt, and Consumer Expectations

This figure displays three-year *Income Growth* (real per capita and from the IRS), *Lagged Debt Overhang* (from Equifax), *Change in Unemployment Rate* (from the BLS), and *Growth of Consumer Expectations* (from the Conference Board) for each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012).




Figure 5: Growth Rate of State Income Inequality, Housing Wealth, and Share of Subprime Borrowers

This figure displays *Income Growth, 90th percentile*, and *Income Growth, 25th percentile* which are the three-year growth rates of real per capita state income for the 90th and 25th percentiles (from the CPS), *Growth of Housing Wealth* (constructed from CoreLogic and Census 2000 data), and *Lagged Subprime Fraction*, which equals the share of individuals in a state whose credit score is lower than 661 (from Equifax).




Figure 6: Change in Mortgage Rates

The figure displays the change in mortgage rates at origination (from Black Knight) for the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). The sample consists of fixed-rate conventional first-lien mortgages originated during each of the subperiods.




Figure 7: Average Value, Economic Significance, and Partial R Squared

The first panel of this figure displays the average value of the variables used in the regressions. We plot the growth rate of the unemployment rate in this panel but we use the change in the unemployment rate in our regressions (and the other panels of this figure). For clearer exposition, we multiplied *Income Growth*, *Growth of 90th Income Percentile* and *Growth of 25th Income Percentile* by 5, *Growth of Consumer Confidence* and *Growth of Share of Cash-Out Refinances* by 10, and we report *Share in Foreclosure in Last Two Years* in percent. The second panel of this figure displays the estimated coefficient for each variable multiplied by one standard deviation (calculated for each time interval in the sample). The third panel of the figure displays the partial R^2 for each variable in each time period. Debt overhang is calculated as total real per capita state-level debt at the beginning of the period (from Equifax).



Table 1: Descriptive Statistics

Variable	2001-2003		2004-2006		2007-2009		2010-2012	
	mean	sd	mean	sd	mean	sd	mean	sd
Growth of Consumption (Retail Sales), %	-0.01	0.11	0.07	0.10	-0.12	0.10	0.10	0.09
Income Growth, %	-5.67	6.44	6.63	8.97	-6.61	7.33	12.32	7.72
Change in Unemployment Rate	1.72	1.14	-1.08	1.02	4.16	1.80	-1.20	1.35
Growth of 90th Income Percentile (State), %	-1.37	4.75	0.34	5.34	2.66	4.32	-2.98	5.13
Growth of 25th Income Percentile (State), %	-2.60	7.38	-3.02	8.26	2.55	6.25	-7.84	6.09
Lagged Share of Income, top 10 Percent, %	30.17	1.61	32.08	2.04	32.63	1.72	31.82	1.66
Growth of Housing Wealth, %	34.79	8.71	35.09	8.69	34.76	8.60	34.60	8.85
Lagged Debt Overhang ($\times \frac{1}{100}$)	11.77	9.03	13.07	11.26	-18.59	10.91	-10.59	6.92
Lagged Subprime Fraction, %	0.27	0.22	0.26	0.22	0.43	0.32	0.29	0.23
Share in Foreclosure in Last Two Years, %	19.55	2.86	19.89	2.83	20.11	2.83	20.17	2.83
Growth of Consumer Confidence (Regional), %	-2.28	0.32	0.38	0.94	-3.25	0.94	1.85	0.36
Growth of Share of Cash-Out Refinancings, %	5.24	3.27	-1.27	3.75	-2.91	2.24	-1.24	1.08

Note: *Growth of Consumption (Retail Sales)* is defined as $\Delta^3 \log(C_{c,t}) = \log(C_{c,t}) - \log(C_{c,t-3})$, the three-year growth rate of consumption proxied by real per capita total county-level retail sales. *Income Growth* is defined similarly using real per capita total county-level gross adjusted income. *Change in Unemployment Rate* is $\Delta^3 UR_{c,t} = UR_{c,t} - UR_{c,t-3}$, the change in unemployment rate over the subperiod. *Growth of 90th (25th) Income Percentile (State)* is the three-year growth rate of 90th (25th) percentiles of real state income distribution. *Lagged Share of Income, top 10 Percent* is the share of real income that belongs to the richest 10 percent of people in a state. *Growth of Housing Wealth* equals (No. of owner occupied housing units $_{c,t} \times$ Median home value $_{c,t}$), which is defined in more details in Section 4. *Lagged Debt Overhang* is real per capita total county debt at the beginning of each subperiod. *Lagged Subprime Fraction* is a fraction of individuals residing in a county with credit scores less than 661 at the beginning of the period. *Share in Foreclosure in Last Two Years* is a share of mortgages in foreclosure relative to all outstanding mortgages in a county over the last two years, measured at the end of each subperiod. *Growth of Consumer Confidence (Regional)* is the three-year growth in consumer confidence index measured at the Census Divisions level. *Growth of Share of Cash-Out Refinancings* is the three-year growth rate of a share of cash-out refinancings relative to the number of all outstanding mortgages. All variables have been winsorised at 2 and 98 percent.

Table 2: Determinants of consumption. Period-by-period regressions.

	2001-2003	2004-2006	2007-2009	2010-2012
Income Growth	0.09** (2.35)	0.08*** (2.94)	0.08** (2.27)	0.10*** (4.04)
Change in Unemployment Rate	-0.01*** (-4.57)	-0.01*** (-5.22)	-0.01*** (-5.56)	-0.01*** (-4.89)
Growth of 90th Income Percentile (State)	0.09* (1.95)	0.10** (2.43)	0.06 (1.23)	-0.09*** (-2.77)
Growth of 25th Income Percentile (State)	0.08** (2.37)	-0.04 (-1.29)	0.05 (1.38)	0.07** (2.36)
Lagged Share of Income, top 10 Percent	0.50*** (3.66)	-0.12 (-1.06)	0.29*** (2.74)	-0.12 (-1.09)
Growth of Housing Wealth	0.04* (1.75)	0.08*** (4.17)	0.08*** (3.48)	0.02 (0.70)
Lagged Debt Overhang	0.39*** (5.39)	-0.46*** (-6.80)	-0.44*** (-6.17)	-0.32*** (-4.77)
Lagged Subprime Fraction	0.00 (0.13)	0.12*** (4.45)	0.03 (0.94)	0.02 (0.84)
Share in Foreclosure in Last Two Years	-2.15* (-1.93)	-2.16** (-2.14)	0.36 (0.45)	-3.03*** (-3.41)
Growth of Consumer Confidence (Regional)	1.67** (2.33)	1.41*** (5.89)	0.43** (2.08)	0.55 (1.07)
Growth of Share of Cash-Out Refinancings	0.04 (0.61)	-0.03 (-0.48)	0.15 (1.59)	0.74*** (3.82)
Constant	-0.01*** (-3.91)	0.07*** (40.61)	-0.12*** (-67.56)	0.10*** (58.62)
Adj. R sq.	0.047	0.127	0.100	0.055
No. obs.	2,770	2,765	2,663	2,662

Note: Cross-sectional regressions over U.S. counties based on the following regression specification for each period:

$$\Delta^3 \log(C_{c,t}) = \alpha + \beta_1 \Delta^3 \log(X_{c,t}) + \beta_2 \widetilde{Y_{c,t-3}} + \beta_3 \Delta^3 \widetilde{UR_{c,t}} + \beta_4 \widetilde{Foreclosure_{c,t}} + \epsilon_{ct},$$

where $\Delta^3 \log(C_{c,t}) = \log(C_{c,t}) - \log(C_{c,t-3})$ is the three-year growth rate of real per capita county-level consumption proxied by total retail sales, $\Delta^3 \log(X_{c,t})$ is a vector of the following variables: *Income Growth* (real per capita total county-level income), *Growth of 90th Income Percentile* and *Growth of 25th Income Percentile* in a given state, *Growth of Housing Wealth* (defined in Section 4), *Growth of Consumer Confidence* in a region, and *Growth of Share of Cash-Out Refinancings* in a county. $Y_{c,t-3}$ are county- or state-level lagged variables: *Lagged Share of Income, top 10 Percent* (state), *Lagged Debt Overhang* (real per capita total county-level debt at the beginning of the period) and *Lagged Subprime Fraction* (the fraction of individuals residing in a county with credit scores less than 661 at the beginning of the period). $\Delta^3 UR_{c,t} = UR_{c,t} - UR_{c,t-3}$ is the change in the unemployment rate over the subperiod, and *Foreclosure_{c,t}* is the share in foreclosure in last two years measured at the end of the period. We demean all independent variables in order for the constant to capture average consumption growth over each three-year interval in the following way: $\widetilde{X}_{c,t} = X_{c,t} - \frac{1}{N} \sum_{c=1}^N X_{c,t}$, for any variable X where c indexes counties and N is the total number of counties in our sample. Consumption, income, debt, and housing wealth are real per capita total aggregates at the county-level. t-statistics based on robust standard errors are reported in parentheses. All variables have been winsorised at 2% and 98%. *** (**) [*] indicate significance at the 1 (5) [10]% level.

Table 3: Partial R^2 and Contribution to R^2

Partial R^2				
	2001-2003	2004-2006	2007-2009	2010-2012
Income Growth	0.28	0.46	0.29	0.66
Change in Unemployment Rate	0.94	1.10	1.41	0.98
Growth of 90th Income Percentile (State)	0.12	0.20	0.06	0.28
Growth of 25th Income Percentile (State)	0.21	0.06	0.07	0.20
Lagged Share of Income, top 10 Percent	0.51	0.05	0.26	0.04
Growth of Housing Wealth	0.11	0.61	0.46	0.02
Growth of Consumer Confidence (Regional)	0.19	1.25	0.13	0.04
Lagged Debt Overhang	0.96	1.54	1.31	0.85
Lagged Subprime Fraction	0.00	0.82	0.04	0.03
Share in Foreclosure in Last Two Years	0.17	0.20	0.01	0.57
Growth of Share of Cash-Out Refinancings	0.02	0.01	0.10	0.76

Contribution to R^2				
	2001-2003	2004-2006	2007-2009	2010-2012
Income Growth	0.06	0.04	0.03	0.12
Change in Unemployment Rate	0.20	0.09	0.14	0.18
Growth of 90th Income Percentile (State)	0.03	0.02	0.01	0.05
Growth of 25th Income Percentile (State)	0.04	0.00	0.01	0.04
Lagged Share of Income, top 10 Percent	0.11	0.00	0.03	0.01
Growth of Housing Wealth	0.02	0.05	0.05	0.00
Growth of Consumer Confidence (Regional)	0.04	0.10	0.01	0.01
Lagged Debt Overhang	0.20	0.12	0.13	0.15
Lagged Subprime Fraction	0.00	0.06	0.00	0.01
Share in Foreclosure in Last Two Years	0.04	0.02	0.00	0.10
Growth of Share of Cash-Out Refinancings	0.00	0.00	0.01	0.14
Not assigned component	0.25	0.50	0.59	0.20

Note: This table shows partial R^2 and contribution to the total R^2 (the partial R^2 as a percent of the total R^2) for each of the explanatory variables used in the baseline regression specification. The partial R^2 s are multiplied by 100. See previous table for variable definitions.

A Appendix

In Figure A.1, we show the three-year consumption growth proxied by growth of total real per capita retail sales, calculated as $\Delta^3 \log(C_{c,t}) = 100 \times (\log(C_{c,t}) - \log(C_{c,t-3}))$ for county c and each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012), using the data sample that has *not been winsorized*. Many of the outliers occur in small counties which may be affected by cross-border shopping or natural disasters. Upon inspecting the data, we found the largest outliers in Gulf Coast counties with large drops in consumption in periods where major hurricanes hit, followed by consumption recoveries during the following periods. We therefore winsorized all variables at 2 and 98 percent for our regressions to make sure our results are not driven by outliers—most results are quite robust to winsorizing but, for example, the effect of income growth varies much less across subperiods with winsorized data.

Our sample is not a perfectly balanced panel. Therefore, when we run our regressions by subperiods we end up with a slightly different number of counties in each subperiod, as shown in Table 2 in the paper. In Table A.1, we demonstrate that our results are robust to estimating our regressions using the sample of counties that have nonmissing variables in all subperiods combined. The results are virtually unchanged.

Figure A.1: County Retail Sales Growth. Not Winsorized.

This figure shows the three-year growth in the real per capita consumption growth, proxied by total retail sales, calculated as $\Delta^3 \log(C_t) = 100 \times (\log(C_t) - \log(C_{t-3}))$ for each of the subperiods: the dot-com recession (2001-2003), the subprime boom (2004-2006), the Great Recession (2007-2009), and the tepid recovery (2010-2012). The data source is Moody's Analytics. The data have *not* been winsorized.




Table A.1: Determinants of consumption. Period-by-period regressions.
The same counties each year.

	2001-2003	2004-2006	2007-2009	2010-2012
Income Growth	0.09** (2.25)	0.08*** (2.96)	0.08** (2.26)	0.10*** (3.90)
Change in Unemployment Rate	-0.01*** (-4.29)	-0.01*** (-4.96)	-0.01*** (-5.41)	-0.01*** (-4.90)
Growth of 90th Income Percentile (State)	0.11** (2.12)	0.04 (1.03)	0.06 (1.37)	-0.09*** (-2.75)
Growth of 25th Income Percentile (State)	0.09*** (2.62)	-0.04 (-1.52)	0.04 (1.24)	0.07** (2.50)
Lagged Share of Income, top 10 Percent	0.46*** (3.33)	-0.13 (-1.14)	0.31*** (2.93)	-0.14 (-1.22)
Growth of Housing Wealth	0.05* (1.95)	0.09*** (4.37)	0.08*** (3.61)	0.02 (0.58)
Lagged Debt Overhang	0.39*** (5.26)	-0.49*** (-7.13)	-0.45*** (-6.24)	-0.30*** (-4.45)
Lagged Subprime Fraction	0.00 (0.10)	0.10*** (3.57)	0.02 (0.80)	0.02 (0.96)
Share in Foreclosure in Last Two Years	-2.40** (-2.11)	-1.94* (-1.86)	0.47 (0.58)	-3.37*** (-3.79)
Growth of Consumer Confidence (Regional)	1.38* (1.86)	1.62*** (6.51)	0.43** (2.07)	0.64 (1.24)
Growth of Share of Cash-Out Refinancings	0.04 (0.53)	-0.04 (-0.66)	0.15 (1.61)	0.73*** (3.72)
Constant	-0.01*** (-3.50)	0.08*** (40.21)	-0.12*** (-67.09)	0.10*** (58.34)
Adj. R sq.	0.045	0.128	0.101	0.055
No. obs.	2,646	2,646	2,646	2,646

Note: See notes to Table 2 *** (**) [*] indicate significance at the 1 (5) [10]% level.