

Dissecting Saving Dynamics: Measuring Wealth, Precautionary, and Credit Effects

March 25, 2012

Christopher Carroll¹ Jiri Slacalek² Martin Sommer³

Abstract

We argue that the U.S. personal saving rate's long stability (from the 1960s through the early 1980s), subsequent steady decline (1980s–2007), and recent substantial increase (2008–2011) can all be interpreted using a parsimonious ‘buffer stock’ model of optimal consumption in the presence of labor income uncertainty and credit constraints. Saving in the model is affected by the gap between ‘target’ and actual wealth, while the target depends on credit conditions and uncertainty. An estimated structural version of the model suggests that increased credit availability accounts for most of the saving rate's long-term decline, while fluctuations in net wealth and uncertainty capture the bulk of the business-cycle variation.

Keywords Consumption, Saving, Wealth, Credit, Uncertainty

JEL codes E21, E32

Web: <http://econ.jhu.edu/people/ccarroll/papers/cssUSSaving/>
PDF: <http://econ.jhu.edu/people/ccarroll/papers/cssUSSaving.pdf>
BibTeX: <http://econ.jhu.edu/people/ccarroll/papers/cssUSSaving-Self.bib>
Slides: <http://econ.jhu.edu/people/ccarroll/papers/cssUSSaving-Slides.pdf>
Archive: <http://econ.jhu.edu/people/ccarroll/papers/cssUSSaving.zip>
(Contains data and estimation software producing paper's results)

¹Carroll: ccarroll@jhu.edu, Department of Economics, 440 Mergenthaler Hall, Johns Hopkins University, Baltimore, MD 21218, <http://econ.jhu.edu/people/ccarroll/>, and National Bureau of Economic Research. ²Slacalek: jiri.slacalek@ecb.europa.eu, European Central Bank, Frankfurt am Main, Germany, <http://www.slacalek.com/>. ³Sommer: msommer@imf.org, International Monetary Fund, Washington, DC, <http://martinsommeronline.googlepages.com/>.

We thank Andra Buca, Kerstin Holzheu and Joanna Slawatyniec for excellent research assistance, and Abdul

Figure 1 Personal Saving Rate in 2007–2011 and Previous Recessions



Notes: The saving rate is expressed as a percent of disposable income. The figure shows the deviation from its value at the start of recession (in percentage points). Historical Range includes all recessions after 1960q1 (when quarterly data become available).

Sources: U.S. Department of Commerce, Bureau of Economic Analysis.

1 Introduction

Fresh interest in the determinants of personal saving has recently been sparked by the remarkable rise in saving during the Great Recession. In the U.S., for example, for the three years after the business cycle peak in 2007 the personal saving rate has remained substantially above its pre-crisis value, and the increase relative to its 2007

Abiad, Mathias Hoffmann, Charles Kramer, María Luengo-Prado, Bartosz Maćkowiak, John Muellbauer, David Robinson, Jonathan Wright and seminar audiences at the Bank of England, the ECB, the 2011 NBER Summer Institute, and the 2012 San Francisco Fed conference on “Structural and Cyclical Elements in Macroeconomics” for insightful comments. We are grateful to John Duca and John Muellbauer for sharing with us their data on credit conditions. The views presented in this paper are those of the authors, and should not be attributed to the International Monetary Fund, its Executive Board, or management, or to the European Central Bank.

value generally exceeded the maximum saving increase after any previous postwar business cycle peak (see Figure 1).¹

Carroll (1992) invoked precautionary motives to explain the tendency of saving to increase during recessions, showing that an older modeling tradition² emphasizing the role of “wealth effects” did not capture cyclical dynamics adequately (particularly for the first of the ‘postmodern’³ recessions in 1990–91 when wealth changed little but saving and unemployment expectations rose markedly).

A largely separate literature has addressed another longstanding puzzle: The steady decline in the U.S. personal saving rate, from over 10 percent of disposable income in the early 1980s to a mere 1 percent in the mid-2000s;^{4,5} here, a prominent theme has been the role of financial liberalization in making it easier for households to borrow.⁶ Some very recent work (Guerrieri and Lorenzoni (2011), Eggertsson and Krugman (2011), Hall (2011)) has argued (though without much attempt at quantification) that a sudden sharp reversal of this credit-loosening trend played a large role in the recent saving rise.⁷

This paper aims to quantify these three channels, both over the long span of historical experience and for the period since the beginning of the Great Recession.

To fix ideas, the paper begins by presenting (in section 2) a tractable ‘buffer stock’ saving model with explicit and transparent roles for each of the influences emphasized above (the precautionary, wealth, and credit channels). The model’s key intuition is that, in the presence of income uncertainty, optimizing households have a target wealth ratio that depends on the usual theoretical considerations (risk aversion, time preference, expected income growth, etc), and on two features that have been harder to incorporate into analytical models: The degree of labor income uncertainty and the availability of credit. Our model yields a tractable analytical solution that can

¹We focus on the U.S. because of its central role in triggering the global economic crisis, and because of the rich existing literature studying U.S. data; but the U.K., Ireland, and many other countries also saw substantial increases in personal saving rates. See Mody, Ohnsorge, and Sandri (2012) for systematic international evidence.

²See Davis and Palumbo (2001) for an exposition, estimation, and review.

³Krugman (2012) seems to have coined the term ‘postmodern’ to capture the change in the pattern of business cycle dynamics dating from the 1990–91 recession (particularly the slowness of employment to recover compared to output). But the pattern has been noted by many other macroeconomists.

⁴We should note here that personal saving rates tend to be one of the most heavily revised data series in the NIPA accounts, and that substantial revisions can occur even many years after the BEA’s “first final” estimates are made (Nakamura and Stark (2007), Deutsche Bank Securities (2012)); the revisions are systematically upward-biased, and often as much as 1-2 percent, so it would not be surprising if some years from now the saving decline in the 2000s appears to be smaller than in the data used here. It seems very unlikely, however, that either the broad trends or the business-cycle frequency movements will be revised greatly; past revisions have tended to be at medium frequencies, and not at either the very low or very high frequencies that tend to provide most of our identification.

⁵Although NIPA accounting conventions impart an inflation-related bias to the measurement of personal saving, the downward trend in saving remains obvious even in an inflation-adjusted measure of the saving rate.

⁶See Parker (2000) for a comprehensive analysis; see Aron, Duca, Muellbauer, Murata, and Murphy (2011) for comparative evidence from the U.S., the U.K., and Japan emphasizing the role of credit conditions in determining saving in all three countries.

⁷A new paper by Challe and Ragot (2012) calibrates a quantitative model with aggregate and idiosyncratic uncertainty and time-varying precautionary saving, and documents that the model can produce a plausible response of consumption to aggregate shocks.

be used to calibrate how much saving should go up in response to an increase in uncertainty, or a negative shock to wealth, or a tightening of liquidity constraints.

We highlight one particularly interesting implication of the model: In response to a permanent worsening in economic circumstances (such as a permanent increase in unemployment risk), consumption initially ‘overshoots’ its ultimate permanent adjustment. This reflects the fact that, when the target level of wealth rises, not only is a higher level of steady-state saving needed to maintain a higher target level of wealth, an immediate *further* boost to saving is necessary to move from the current (inadequate) level of wealth up to the new (higher) target. An interesting implication is that if the economy suffers from adjustment costs for overall aggregate demand (as macroeconomic models strongly suggest), an optimizing government might wish to counteract the component of the consumption decline that reflects ‘overshooting.’ In an economy rendered non-Ricardian by liquidity constraints and/or uncertainty, this provides a potential rationale for countercyclical fiscal policy, either targeted at households or to boost components of aggregate demand other than household spending in order to offset the temporary downward overshooting of consumption.

After section 3’s discussion of data and measurement issues, section 4 presents a reduced-form empirical model, motivated by the theory, that attempts to measure the relative importance of each of these effects (precautionary, wealth, and credit) for the U.S. personal saving rate. An OLS regression of the personal saving rate on proxies for the model’s three variables finds a statistically significant and economically important role for all three. The model’s estimated coefficients imply that the largest contributor to the decline in consumption during the Great Recession was the collapse in household wealth, with the increase in precautionary saving also making a substantial contribution; the role of measured changes in credit availability is estimated to have played a substantially smaller (though not negligible) role.

Section 5 constructs a more explicit relationship between the theoretical model and the empirical results, by making a direct identification between the model’s parameters (like unemployment risk) and the corresponding empirical objects (like households’ unemployment expectations constructed using the Thomson Reuters/University of Michigan’s Surveys of Consumers). We show that the structural model fits the data essentially as well as the reduced form model, but with the usual advantage of structural models that it is possible to use the estimated model to provide a disciplined investigation of quantitative theoretical issues such as whether there is an interaction between the precautionary motive and credit constraints. (We find some evidence that there is).

2 Theory: Target Wealth and Credit Conditions

Carroll and Toche (2009) (henceforth CT) provide a tractable framework for analyzing

the impact of nonfinancial uncertainty, in the specific form of unemployment risk, on optimal household saving. Carroll and Jeanne (2009) show that the lessons from the individual’s problem, solved below, carry over with little modification to the characterization of the behavior of aggregate variables in a small open economy. A satisfactory closed-economy general equilibrium analysis remains elusive (though see Challe and Ragot (2012) for a valiant effort.) Such an effort would be useful because a crucial question is the extent to which each of the influences we measure is an “impulse” versus the extent to which it is a “propagation mechanism.” (In the Great Recession, the collapse in consumer confidence seems to have preceded the credit tightening; the wealth decline began before either of the other two variables moved, but its sharpest contractions came after both other variables had deteriorated sharply).

The consumer maximizes the discounted sum of utility from an intertemporally separable CRRA utility function $u(\bullet) = \bullet^{1-\rho}/(1-\rho)$ subject to the dynamic budget constraint:

$$m_{t+1} = (m_t - c_t)R + \ell_{t+1}W_{t+1}\xi_{t+1},$$

where next period’s market resources m_{t+1} are the sum of current market resources net of consumption c_t , augmented by the (constant) interest factor $R = 1 + r$, and with the addition of labor income. The level of labor income is determined by the individual’s productivity ℓ (lower case letters designate individual-level variables), the (upper-case) aggregate wage W_{t+1} (per unit of productivity) and a zero–one indicator of the consumer’s employment status ξ .

The assumption that makes the model tractable is that unemployment risk takes a particularly stark form: Employed consumers face a constant probability \mathcal{U} of becoming unemployed, and, once unemployed, the consumer can never become employed again.⁸ Under these assumptions, CT derive a formula for the steady-state target \tilde{m} that depends on unemployment risk \mathcal{U} , the interest rate r , the growth rate of wages ΔW , relative risk aversion ρ , and the discount factor β :⁹

$$\tilde{m} = f(\underbrace{\mathcal{U}}_{(+)}, \underbrace{r}_{(+)}, \underbrace{\Delta W}_{(-)}, \underbrace{\rho}_{(+)}, \underbrace{\beta}_{(+)}) \tag{1}$$

Target m increases with unemployment risk, because in response to higher uncer-

⁸Of course, if a starting population of such consumers were not refreshed by an inflow of new employed consumers, the population unemployment rate would asymptote to 100 percent. This problem can easily be addressed by introducing explicit demographics (which do not affect the optimization problem of the employed): Each period new employed consumers are born and a fraction of existing households dies, as in Carroll and Jeanne (2009). Because demographic effects are very gradual, the implications of the more complicated model are well captured by the simpler model presented here that ignores demographics and the behavior of the unemployed population.

⁹Specifically, the steady-state target wealth can be approximated as

$$\tilde{m} = 1 + \frac{1}{\mathfrak{p}_r(\hat{\mathfrak{p}}_\gamma/\mathcal{U}) - \mathfrak{p}_\gamma},$$

where $\mathfrak{p}_r = \log((R\beta)^{1/\rho})/R$, $\mathfrak{p}_\gamma = \log((R\beta)^{1/\rho})/\Gamma$, $\hat{\mathfrak{p}}_\gamma = \mathfrak{p}_\gamma(1 + \mathfrak{p}_\gamma\omega/\mathcal{U})$, $\Gamma = (1 + \Delta W)/(1 - \mathcal{U})$ and $\omega = (\rho - 1)/2$.




Figure 2 Consumption Function (Stable Arm of Phase Diagram)

tainty, consumers choose to build up a larger precautionary buffer of wealth to protect their spending. (The increase in \mathcal{U} is a pure increase in risk (a mean-preserving spread in human wealth) because productivity is assumed to grow by the factor $1/(1 - \mathcal{U})$ each period, $l_{t+1} = l_t/(1 - \mathcal{U})$ (see Carroll and Toche (2009), p. 6)). A higher interest rate increases the rewards to holding wealth and thus increases the amount held. Faster income growth translates into a lower wealth target because households who anticipate higher future income consume more now in anticipation of their future prosperity (the ‘human wealth effect’). Finally, risk aversion and the discount factor have effects on target wealth that are qualitatively similar to the effects of uncertainty and the interest rate, respectively. While the unemployment risk in Carroll and Toche (2009) is of a simple form, the key mechanisms at work are the same as those in more sophisticated setups with a realistic specification of uninsurable risks (building on the work of Bewley (1977), Skinner (1988), Zeldes (1989), Deaton (1991), Carroll (1992), Carroll (1997) and many others).

Figure 2 shows the phase diagram for the CT model. The consumption function is indicated by the thick solid locus, which is the saddle path that leads to the steady state at which the ratios of both consumption and market resources to income (c and m) are constant.¹⁰

This consumption function can be used directly to analyze the consequences of

¹⁰For a detailed intuitive exposition of the model, see <http://econ.jhu.edu/people/ccarroll/public/lecturenotes/consumption/tractablebufferstock/>.




Figure 3 A Wealth Shock

an exogenous shock to wealth of the kind contemplated in the old “wealth effects” literature, or in the AEA Presidential Address of Hall (2011).¹¹ The consequences of a pure shock to wealth are depicted in figure 3 and are straightforward: Consumption declines upon impact, to a level below the value that would leave m^e constant (the leftmost red dot); because consumption is below income, m^e (and thus c^e) rises over time back toward the original target (the sequence of dots).

The model solved in CT deliberately omitted explicit liquidity constraints in order to emphasize the point that uncertainty induces concavity of the consumption function (that is, a higher marginal propensity to consume for people with low levels of wealth) even in the absence of constraints (for a general proof of this proposition, see Carroll and Kimball (1996)). Indeed, because the employed consumer is always at risk of a transition into the unemployed state where income will be zero, the ‘natural borrowing constraint’ in this model prevents the consumer from ever choosing to go into debt, because an indebted unemployed consumer with zero income might be forced to consume zero or a negative amount (incurring negative infinity utility) in order to satisfy the budget constraint.

¹¹Like that literature, we take the wealth shock to be exogenous. It is clear from the prior literature starting with Merton (1969) and Samuelson (1969) that not much would change if a risky return were incorporated and the wealth shock were interpreted as a particularly bad realization of the stochastic return on assets. The much more difficult problem of constructing a plausible general equilibrium theory of endogenous asset pricing that could justify the observed wealth shocks has not yet been satisfactorily solved, which is why we follow Hall (2011) in treating wealth shocks at the beginning of the Great Recession as exogenous.

We make only one modification to the CT model for the purpose at hand: We introduce an ‘unemployment insurance’ system that guarantees a positive level of income for unemployed households. In the presence of such insurance, households with low levels of market resources will be willing to borrow because they will not starve even if they become unemployed. This change induces a leftward shift in the consumption function by an amount corresponding to the present discounted value of the unemployment benefit. The consumer will limit his indebtedness, however, to an amount small enough to guarantee that consumption will remain strictly positive even when unemployed (this requirement defines the ‘natural borrowing constraint’ in this model).

We could easily add a tighter ‘artificial’ liquidity constraint, imposed exogenously by the financial system, that would prevent the consumer from borrowing as much as the natural borrowing constraint permits. But Carroll (2001) shows that the effects of tightening an artificial constraint are qualitatively and quantitatively similar to the effects of tightening the natural borrowing constraint; while we do not doubt that artificial borrowing constraints exist and are important, we do not incorporate them into our framework since we can capture their consequences by manipulating the natural borrowing constraint that is already an essential element of the model. Indeed, using this strategy, our empirical estimates below will interpret the process of financial liberalization which began in the U.S. in the early 1980s and arguably continued until the eve of the Great Recession as the major explanation for the long downtrend in the saving rate.

Figure 4 shows that the model reproduces the standard result from the existing literature (see, e.g., Carroll (2001), Muellbauer (2007), Guerrieri and Lorenzoni (2011), Hall (2011)): Relaxation of the borrowing constraint (from an initial position of 0. in which no borrowing occurs, to a new value in which the natural borrowing limit is \underline{h} implying minimum net worth of $-\underline{h}$) leads to an immediate increase in consumption for a given level of resources. But over time, the higher spending causes the consumer’s level of wealth to decline, forcing a corresponding gradual decline in consumption until wealth eventually settles at its new, lower target level. (For vivid illustration, parameter values for this figure were chosen such that the new target level of wealth is negative; that is, the consumer would be in debt, in equilibrium).

Rather than presenting yet another variant of the phase diagram, we instead illustrate our next experiment by showing the dynamics of the saving rate rather than the level of consumption over time. (Since both saving and consumption are strictly monotonic functions of m^e , there is a mathematical equivalence between the two ways of presenting the results).

Figure 5 shows the consequences of a permanent increase in unemployment risk \mathcal{U} : An immediate jump in the saving rate, followed by a gradual decline toward a new equilibrium rate that is higher than the original one.

Qualitatively, the effects of an increase in risk are essentially the opposite of a credit




Figure 4 Relaxation of a Natural Borrowing Constraint from 0. to \underline{h}




Figure 5 Dynamics of the Saving Rate after an Increase in Unemployment Risk

loosening: In response to a human-wealth-preserving spread in unemployment risk, the level of consumption falls sharply as consumers begin the process of accumulation toward a higher target wealth ratio.¹² The figure illustrates the ‘overshooting’ proposition mentioned in the introduction: All of the initial increase in saving reflects a drop in consumption (by construction, the mean-preserving spread in unemployment risk leaves current income unchanged), and consumption recovers only gradually toward its ultimately higher target. For a long time, the saving rate remains above either its pre-shock level or its new target.

Economists’ instinct (developed in complete-markets and perfect-foresight models) is that privately optimal behavior also usually has some plausible claim to reflect a socially efficient outcome. This is emphatically *not* the case for movements in precautionary saving against idiosyncratic risk in models with imperfect capital markets. It has long been known that such precautionary saving generates socially ‘excessive’ saving (see, e.g., Aiyagari (1993)). So the presumption from economic theory is that the increase in the precautionary motive following an increase in uninsurable risk is socially inefficient. The inefficiency would be even greater if we were to add to our model a production sector like the one that has become standard in DSGE models in which there are costs of adjustment to the amount of aggregate investment (e.g., Christiano, Eichenbaum, and Evans (2005)).

While the implications for optimal fiscal policy are beyond the scope of our analysis, it is clear that a number of policies could either mitigate the consumption decline (e.g., an increase in social insurance) or replace the corresponding deficiency in aggregate demand (e.g., by an increase in government spending). We leave further exploration of these ideas to later work, or other authors.

One objection to the model might be that its extreme assumption about the nature of unemployment risk (once unemployed, the consumer can never become reemployed) calls into question its practical usefulness except as a convenient stylized treatment of the logic of precautionary saving. Our view is that such a criticism would be misplaced, for several reasons. First, when unemployment risk in the model is set to zero, it collapses to the standard Ramsey model that has been a workhorse for much of macroeconomic analysis for the past 40 years (see Carroll and Toche (2009) for details). It seems perverse to criticize the model for moving at least a step in the direction of realism by introducing a precautionary motive into that framework. Second, this paper’s authors have been active participants in the literature that builds far more realistic models of precautionary saving, but our considered judgment is that in the present context the virtues of transparency and simplicity far outweigh the model’s cost in realism. Models are metaphors, not high-definition photographs,

¹²The model is specified in such a way that an increase in the parameter \bar{U} that we are calling the ‘unemployment risk’ here actually induces an offsetting increase in the expected mean level of income (an increase in \bar{U} is a mean-preserving spread in the relevant sense); the spending of a consumer with certainty-equivalent preferences therefore would not change in response to a change in \bar{U} , so we can attribute all of the increase in the saving rate depicted in the figure to the precautionary motive.

and if a certain flexibility of interpretation is granted to use a simple model that has most of the right parts, more progress can sometimes be made than by building a state-of-the-art Titanic.

In sum, the model emphasizes three factors that affect saving and that might vary substantially over time. First, because the precautionary motive diminishes as wealth rises, the saving rate is a declining function of market resources m_t . Second, since an expansion in the availability of credit reduces the target level of wealth, looser credit conditions (designated CEA_t , for reasons articulated below) lead to lower saving. Finally, higher unemployment risk \mathcal{U}_t results in greater saving for precautionary reasons.

The framework thus suggests that if proxies for these variables can be found, a reduced-form regression for the saving rate s_t

$$s_t = \gamma_0 + \gamma_m m_t + \gamma_{CEA} CEA_t + \gamma_{\mathcal{U}} \mathcal{U}_t + \gamma' X_t + \varepsilon_t \quad (2)$$

should satisfy the following conditions:

$$\gamma_m < 0, \quad \gamma_{CEA} < 0, \quad \gamma_{\mathcal{U}} > 0, \quad (3)$$

where CEA_t denotes the ‘‘Credit Easing Accumulated’’ index, a measure of credit supply (described in detail below), and the vector X_t collects other drivers of saving that are outside the scope of the model, such as demographics, corporate and government saving, etc. We estimate regressions of the form (2) in section 4 below.

To economists steeped in the wisdom of Irving Fisher (1930) according to whom the consumption path is determined by lifetime resources independently of the income path (‘Fisherian separation holds’), equation (2) may seem like a throwback to the bad old days of nonstructural Keynesian estimation of the kind that fell into disrepute after spectacular failures in the 1970s. Below, however, we will show that, at least under our assumptions, a reduced form estimation of such an equation can in principle yield estimates of ‘‘structural’’ parameters like the time preference rate. (An important part of the reason this exercise is not implausible is that, with the exception of a few easily identified episodes, the growth rate of personal income is not very far from a random walk with drift, justifying the identification of actual aggregate personal income with ‘permanent income’ in a Friedmanian sense).¹³

¹³More precisely, an empirical decomposition of NIPA personal disposable income into permanent and transitory components (in which income consists of unobserved random walk with drift and white noise) assigns almost all variation in (measured) income to its permanent component, so that a ratio to actual income will coincide almost perfectly with a ratio to estimated ‘permanent’ income. This is not surprising because, as is well-known (and also documented in Appendix 2), it is difficult to reject the proposition that almost all shocks to the level of aggregate income are permanent; autocorrelation functions and partial autocorrelation functions indicate that log-level of disposable income is close to a random walk; see our further discussion in Appendix 2.

3 Data and Measurement Issues

Before presenting estimation results we introduce our dataset. Because our empirical measure of credit conditions begins in 1966q2, our analysis begins at that date and extends (at the present writing) through 2011q1.^{14,15} The saving rate is from the BEA’s National Income and Product Accounts and is expressed as a percentage of disposable income.^{16,17}

One objection to our analysis might be that some items that are included in personal consumption expenditures (in particular, spending on highly durable goods like automobiles) are more properly treated as saving in a nonfinancial form rather than spending. We acknowledge this point, but our view is that its force is easily exaggerated, for several reasons. Perhaps the most important is that in the short run, the most urgent purpose for modeling of this kind is to provide guidance to policymakers who need to assess the likely path of consumer expenditures as defined in the national accounts, since such direct spending is what contributes to GDP and can be influenced by both fiscal and monetary policy. Policymakers who were offered a model that fitted ‘consumption’ as abstractly defined by theory, but did not say much about NIPA PCE, would probably prefer our model. A second response is that, with respect to long-run trends, the depreciation rate of durable goods included in NIPA PCE (even including automobiles) is high enough that we would not expect much bias from the ‘durables’ problem over the course of a 45 year estimation period. A further point is that the theoretical forces in which we are most interested, the roles of liquidity constraints and precautionary saving, are precisely those forces that have been shown most seriously to undermine the implications of the intellectual framework (perfect foresight, perfect capital markets, no adjustment costs, etc) that justifies the treatment of automobile purchases as equivalent to saving in a drivable investment vehicle. Finally, any attempt to perform an analysis similar to ours but using a more theoretically “pure” measure of consumption quickly descends into a

¹⁴Most time series were downloaded from Haver Analytics, and were originally compiled by the Bureau of Economic Analysis, the Bureau of Labor Statistics or the Federal Reserve.

¹⁵We are reluctant to use more recent data because personal saving rate statistics are subject to particularly large revisions until the data have been subjected to at least one annual revision (Deutsche Bank Securities (2012); Nakamura and Stark (2007)).

¹⁶As a robustness check, we have also re-estimated our models with alternative measures of saving: Gross household saving as a fraction of disposable income, gross and net private saving as a fraction of GDP, inflation-adjusted personal saving rate and two measures of saving from the Flow of Funds (with/without consumer durables). The inflation-adjusted saving rate deducts from saving the erosion in the value of money-denominated assets due to inflation. The Flow of Funds (FoF) calculates saving as the sum of the net acquisition of financial assets and tangible assets minus the net increase in liabilities. Because this FoF-based measure is substantially more volatile, the fit of the model is worse than for the NIPA-based PSR. However, the main messages of the paper remain unchanged.

¹⁷Many reasonable objections can be made to this, or any other, specific measure of the personal saving rate, including the treatment of durable goods, the treatment of capital gains and losses, and so on. While some defense of the NIPA measure could be made in response to many of these challenges, such defenses would take us too far afield, and we refer the reader to the extensive discussions of these measurement issues that date at least back to Friedman (1957).

morass of arbitrary judgments like whether spending on holiday travel is “durable” because memories can last a lifetime. While some existing papers have made a stab at drawing such lines (e.g., Blinder and Deaton (1985)), our view is that at the quarterly frequency probably the principal categories of goods that are pretty clearly nondurable are fresh fruits and vegetables. (Canned ones can last much longer than a quarter, and even meat and fish can be frozen).

Market resources m_t are measured as 1 plus the ratio of household net worth to disposable income, in line with the model.¹⁸ Our measure of credit supply conditions, which we call the Credit Easing Accumulated index (CEA, see Figure 6), is constructed in the spirit of Muellbauer (2007) and Duca, Muellbauer, and Murphy (2010) using the question on consumer installment loans from the Federal Reserve’s Senior Loan Officer Opinion Survey (SLOOS) on Bank Lending Practices (see also Fernandez-Corugedo and Muellbauer (2006) and Hall (2011)). The question asks about banks’ willingness to make consumer installment loans now as opposed to three months ago (we use this index because it is available since 1966; other measures of credit availability, such as for mortgage lending, move closely with the index on consumer installment loans over the sample period when both are available). To calculate a proxy for the *level* of credit conditions, the scores from the survey were accumulated, weighting the responses by the debt–income ratio to account for the increasing trend in that variable.¹⁹ (The index is normalized between 0 and 1 to make the interpretation of regression coefficients straightforward.)


The CEA index is taken to measure the availability/supply of credit to a typical household through factors other than the level of interest rates—for example, through loan to value and loan to income ratios, availability of mortgage equity withdrawal and mortgage refinancing. The broad trends in the CEA index correlate strongly with measures financial reforms of Abiad, Detragiache, and Tressel (2008), and measures of banking deregulation of Demyanyk, Ostergaard, and Rensen (2007) (see panel A of their Figure 1, p. 2786).²⁰ In addition, they seem to reflect well the key developments of the U.S. financial market institutions as described in McCarthy and Peach (2002), Dynan, Elmendorf, and Sichel (2006), Green and Wachter (2007), Campbell and Hercowitz (2009), and Aron, Duca, Muellbauer, Murata, and Murphy (2011), among others, which we summarize as follows. Until the early 1980s, the U.S. consumer lending markets were heavily regulated and segmented. After the phaseout of interest rate controls beginning in the early 1980s, the markets became more competitive,

¹⁸This variable is lagged by one quarter to account for the fact that data on net worth are reported as the end-of-period values.

¹⁹As in Muellbauer (2007), we use the question on consumer installment loans rather than mortgages because the latter is only available starting in 1990q2 and the question changed in 2007q2. Our CEA index differs from Muellbauer (2007)’s Credit Conditions Index in that Muellbauer accumulates raw answers, not weighting them by the debt–income ratio.

²⁰Duca, Muellbauer, and Murphy (2011) document an increasing trend in loan to value ratios for first-time home buyers (in data from the American Housing Survey, 1979–2007), an indicator which is arguably to some extent affected by fluctuations in demand.

Figure 6 The Credit Easing Accumulated (CEA) Index



Notes: Shading—NBER recessions.

Sources: Federal Reserve, accumulated scores from the question on change in the banks' willingness to provide consumer installment loans from the Senior Loan Officer Opinion Survey on Bank Lending Practices, <http://www.federalreserve.gov/boarddocs/snloansurvey/>.

spurring financial innovations that led to greater access to credit. Technological progress leading to new financial instruments and better credit screening methods, a greater role of nonbanking financial institutions, and the increased use of securitization all contributed to the dramatic rise in credit availability from the early 1980s until the onset of the Great Recession in 2007. The subsequent significant drop in the CEA index was associated with the funding difficulties and de-leveraging of financial institutions. As a caveat, it is important to acknowledge that CEA might to some degree be influenced by developments from the demand rather than the supply side of the credit market. But whatever its flaws in this regard, indexes of this sort seem to be gaining increasing acceptance as the best available measures of credit supply (as distinguished from credit demand).²¹

We measure a proxy $\mathbb{E}_t u_{t+4}$ for unemployment risk \mathcal{U}_t using re-scaled answers to the question about the expected change in unemployment in the Thomson Reuters/University of Michigan Surveys of Consumers.²² In particular, we estimate $\mathbb{E}_t u_{t+4}$ using fitted values $\Delta_4 \hat{u}_{t+4}$ from the regression of the four-quarter-ahead change in unemployment rate $\Delta_4 u_{t+4} \equiv u_{t+4} - u_t$ on the answer in the survey, summarized with a balance statistic UExp_t^{BS} :

$$\begin{aligned}\Delta_4 u_{t+4} &= \alpha_0 + \alpha_1 \text{UExp}_t^{BS} + \varepsilon_{t+4}, \\ \mathbb{E}_t u_{t+4} &= u_t + \Delta_4 \hat{u}_{t+4}.\end{aligned}$$

The coefficient α_1 is highly statistically significant (indicating that households do have substantial information about the direction of future changes in the unemployment rate). Our $\mathbb{E}_t u_{t+4}$ series, which—as expected—correlates strongly with unemployment rate and precedes its dynamics, is shown in Figure 7.

4 Reduced-Form Saving Regressions

Before proceeding to structural estimation of the model of section 2 we investigate a simple reduced-form benchmark:


$$s_t = \gamma_1 + \gamma_m m_t + \gamma_{\text{CEA}} \text{CEA}_t + \gamma_{Eu} \mathbb{E}_t u_{t+4} + \gamma_t t + \gamma' X_t + \varepsilon_t. \quad (4)$$

Such a specification can be readily estimated using OLS or IV estimators, and at a minimum can be interpreted as summarizing basic stylized facts about the data.

²¹We have verified that our results do not materially change when we use the credit conditions index of Duca, Muellbauer, and Murphy (2010), which differs from our CEA in that Duca, Muellbauer, and Murphy explicitly remove identifiable effects of interest rates and the macroeconomic outlook from the SLOOS data using regression techniques. Since the results are similar using both measures, our interpretation is that our measure is at least not merely capturing the most obvious cyclical components of credit demand. As reported below, our results also do not change when we use the Financial Liberalization Index of Abiad, Detragiache, and Tressel (2008)—which is based on the readings of financial laws and regulations—as an instrument for CEA.

²²The relevant question is: “How about people out of work during the coming 12 months—do you think that there will be more unemployment than now, about the same, or less?”


Figure 7 Unemployment Risk $\mathbb{E}_t u_{t+4}$ and Unemployment Rate (Percent)



Legend: Unemployment rate: Thin black line, Unemployment risk $\mathbb{E}_t u_{t+4}$: Thick red/grey line. Shading—NBER recessions.

Sources: Thomson Reuters/University of Michigan Surveys of Consumers, <http://www.sca.isr.umich.edu/main.php>, Bureau of Labor Statistics.

Figure 8 The Fit of the Baseline Model and the Time Trend—Actual and Fitted PSR (Percent of Disposable Income)



Legend: Actual PSR: Thin black line, Baseline model: Thick red/grey line, Time trend: Dashed black line. Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors' calculations.

Table 1 reports the estimated coefficients from several variations on equation (4). The first four columns show univariate specifications in which the saving rate is in turn regressed on each of the three determinants analyzed above: Wealth, credit conditions, and unemployment risk. In each specification we include the time trend to investigate how much each regressor contributes to explaining the PSR beyond the portion that can be captured mechanically by a linear time effect. The three coefficients have the signs predicted by the model of section 2 and are statistically significant. Univariate regressions capture up to 85 percent of variation in saving.

But the univariate models on their own do not adequately describe the dynamics of the PSR. As the model labeled “All 3” in the fifth column shows, the three key variables of interest—wealth and credit conditions—jointly explain roughly 90 percent of the variation in the saving rate over the past five decades. As expected, the point estimates again indicate a strong negative correlation between saving and net wealth and credit conditions and a positive correlation with unemployment risk. Interestingly, once the three variables are included jointly, the time trend ceases to be significant, which is in line with the fact that the three models in columns 2–4 have higher \bar{R}^2 than the univariate model with the time trend only.²³

The specification in column 5 (All 3) suggests that a more parsimonious version of the model without the time trend reported in column 6 (Baseline)—and also suggested by the structure of section 2—neatly summarizes the key features of the saving rate. The estimated coefficient on net wealth implies the (direct) long-run marginal propensity to consume of about 1.2 cents out of a dollar of (total) wealth. The value is low compared to much of the literature, which typically estimates a marginal propensity to consume out of wealth (MPCW) of about 3–7 cents without explicitly accounting for credit conditions.²⁴ However, a univariate model regressing the PSR just on net wealth (not reported here), implies an MPCW of 4.3 percent. These results suggest that much of what has been interpreted as pure “wealth effects” in the prior literature may actually have reflected precautionary or credit availability effects that are correlated with wealth.

The coefficient on the Credit Easing Accumulated index is highly statistically significant with a t statistic of -10.7 . The point estimate of γ_{CEA} implies that increased access to credit during the sample period ending in 2007 (before the Great Recession) reduced the PSR by about 6 percentage points of disposable income. In the aftermath of the Recession, the CEA index declined between 2007 and 2010 by

²³Estimating univariate saving regressions without the time trend, results in higher \bar{R}^2 for wealth and the credit conditions—0.72 and 0.80, respectively—that for the “time” model in column one (0.70). (Because unemployment risk is not trending, it captures relatively little variation in saving on its own (about 10 percent) but is important *in addition* to the two other factors, as illustrated in columns 4 and 5.)

²⁴See, for example, Skinner (1996), Ludvigson and Steindel (1999), Lettau and Ludvigson (2004), Case, Quigley, and Shiller (2005), and Carroll, Otsuka, and Slacalek (2011). See Muellbauer (2007) and Duca, Muellbauer, and Murphy (2010) for a model which includes a measure of credit conditions in the consumption function.

roughly 0.11 as credit supply tightened, contributing roughly 0.64 percentage point to the increase in the PSR (see the discussion of Table 3 below for more detail).

Figure 8 further illustrates why we find the “baseline” specification in column 6 more appealing than the more atheoretical model with a linear time trend. The trends in saving and the CEA are both non-linear, moving consistently with each other even within our sample and often persistently departing from the linear trend (as indicated by the time-only model’s substantially lower \bar{R}^2). In addition, it is likely that the time-only model will become increasingly problematic as observations beyond our sample accumulate, arguably providing additional evidence on the structural break in the time model during the Great Recession.²⁵

Finally, the last model investigates the joint effect of credit conditions and unemployment risk. The structural model of section 2 implies that uncertainty affects saving more strongly when credit constraints bind tightly; the model in column 7 (Interact) confirms the prediction with a (borderline) significant negative interaction term between the CEA and unemployment risk.²⁶

Table 2 presents a second battery of specification checks of the baseline model (shown again for reference in the first column). The second model (Uncertainty) investigates the effects of adding to the baseline regression an alternative proxy for uncertainty: the Bloom, Floetotto, and Jaimovich (2009) index of macroeconomic and financial uncertainty.²⁷ The new variable is statistically insignificant and the coefficients on the previously included variables are broadly unchanged, suggesting that our baseline uncertainty measure is more appropriate for our purposes (which makes sense, as personal saving is conducted by persons, whose uncertainty is likely better captured by our measure of labor income uncertainty than by the Bloom, Floetotto, and Jaimovich (2009) measure of firm-level shocks).

The third model (Lagged s_{t-1}) explores the implications of adding lagged saving to the list of regressors. Often in empirical macroeconomics, the addition of the lagged dependent variable is unjustified by the underlying theory, but nevertheless is required for the model to fit the data. Here, however, serial correlation in saving is a direct implication of the model (below we will show that the degree of serial correlation implied by the model matches the empirical estimate fairly well). The implication arises because deviations of actual wealth from target wealth ought to be long-lasting if the saving rate cannot quickly move actual wealth to the target. As expected, the coefficient is highly statistically significant. However, this positive autocorrelation only captures near-term stickiness and has little effect on the long-run dynamics of saving. Indeed, the coefficients from the baseline roughly equal their

²⁵Reliable PSR data only start in 1959 and document that the downward trend in saving started around 1975, so that our sample is actually quite favorable to the time-only model; it would have considerably more difficulty with a sample that included 10 pre-sample years without a discernable trend.

²⁶Adding an interaction term between the CEA and wealth results in a borderline significant positive estimate, which is in line with the concavity of the consumption function, show in Figure 2.

²⁷See Baker, Bloom, and Davis (2011) for related work measuring *economic policy* uncertainty.

long-term counterparts from the model with lagged saving rates (that is, coefficient estimates pre-multiplied by 2.5, or $1/(1 - \gamma_s) = 1/(1 - 0.60)$).²⁸

The fourth model (Debt) explores the role of the debt–income ratio. The variable could be relevant for two reasons. First, it could partly account for the fact that debt is held by a different group of people than assets and consequently net worth might be an insufficient measure of wealth. Second, debt might also reflect credit conditions (although—as mentioned above—we prefer the CEA index because in principle it isolates the role of credit supply from demand). The regression can thus also be interpreted as a horse-race between the CEA and the debt–income ratio. In any case, while the coefficient γ_d has the correct (negative) sign, it is statistically insignificant and its inclusion does not substantially affect estimates obtained under the baseline specification.

The fifth model (Full Controls) controls for the effects of other potential determinants of household saving: expected real interest rates, expected income growth, and government and corporate saving (both measured as a percent of GDP).²⁹ Some of these factors are statistically significant, but all are inconsequential in economic terms. Figure 9 makes it clear that while these additional factors were potentially important during specific episodes (especially in the early 1980s), they have on average had only a limited impact on U.S. household saving. The negative coefficient on corporate saving is consistent with the proposition that households may ‘pierce the corporate veil’ to some extent³⁰ but there is no evidence for any interaction between personal and government saving. One interpretation of this is that ‘Ricardian’ effects that some prior researchers have claimed to find might instead reflect reverse causality: Recessions cause government saving to decline at the same time that personal saving increases (high unemployment, falling wealth, restricted credit) but for reasons independent of the Ricardian logic (reduced tax revenues and increased spending on automatic stabilizers, e.g.). Since we are controlling directly for the variables (wealth, unemployment risk, credit availability) that were (in this interpretation) proxied by government saving, we no longer find any effect of government saving on personal saving.


When the model is estimated only using the post-1980 data in the sixth column (Post-1980), its fit measured by the \bar{R}^2 actually improves, in contrast with many other economic relationships, whose goodness-of-fit deteriorated in the past 20 years. The F test is consistent with the proposition that the coefficients of the regression have not changed over the sample.

²⁸Note that with the inclusion of lagged saving, the Durbin–Watson statistic becomes close to 2, suggesting that whatever serial correlation exists in the other specifications reflect simple first order autocorrelation of the errors.

²⁹Expected real interest rates and expected income growth are constructed using data from the Survey of Professional Forecasters of the Philadelphia Fed.

³⁰Regressions with total private saving as a dependent variable yield qualitatively similar results as our baseline estimates in Table 1.

Figure 9 The Fit of the Baseline Model and the Model with Full Controls (of Table 2)—Actual and Fitted PSR (Percent of Disposable Income)



Legend: Actual PSR: Thin black line, Baseline model: Thick red/grey line, Model with full control variables: Dashed black line. Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors' calculations.

Finally, to explore how much endogeneity may matter,³¹ the specification “IV” re-estimates the baseline specification using the IV estimator. Instruments are the lags of net wealth, unemployment risk and—crucially—the Financial Liberalization Index of Abiad, Detragiache, and Tressel (2008) (described in Appendix 1). The FLI is an alternative measure of credit conditions constructed using the records about legal and regulatory changes in the banking sector. The index intends to capture exogenous changes in credit conditions. While it is a rough approximation as it reflects only the most important events (see also Figure 14 in Appendix 1), the profile of the FLI matches well that of the CEA. The estimated coefficients remain broadly unchanged compared with the baseline specification.

We have also estimated specifications with other variables, whose detailed results we do not report. As in Parker (2000), demographic variables, like the old-age dependency ratio, were insignificant in our regressions. The importance of population aging in cross-country studies of household saving (for example, Bloom, Canning, Mansfield, and Moore (2007) and Bosworth and Chodorow-Reich (2007)) appears to be largely driven by the experience of Japan and Korea—countries well ahead of the United States in the population aging process.

To address a potential criticism that saving rate regressions are difficult to interpret because aggregate income shocks reflect a mix of transitory and persistent factors, we have also re-estimated our regressions with alternative measures of disposable income (see Appendix 2) which exclude a range of identifiable temporary shocks such as fiscal stimulus and extreme weather. There was little econometric evidence that transitory movements in aggregate disposable income are substantial and our econometric results basically did not change.³²

Table 3 reports in-sample fit of the baseline model and the model Interact with the CEA–uncertainty interaction term of Table 1, and the contributions of the individual variables to the explained increase in the saving rate between 2007 and 2010. Two principal conclusions emerge. First, both models (especially the latter) are able to capture well the observed change in the saving rate. Second, the key explanatory factors in saving were the changes in wealth and uncertainty, with credit conditions (as measured by CEA) playing a less important role. While the change in the trajectory of the CEA index is quite striking (see Figure 6), and may explain the sudden academic interest in the role of household credit over the business cycle (see the papers cited in the introduction), this evidence suggests that the rise in saving cannot be mainly attributed to the decline in credit availability. If correct, this finding is particularly important at the present juncture because it suggests that however much the health

³¹As mentioned above, wealth is lagged by one quarter to alleviate endogeneity in OLS regressions. However, a standard concern about reduced-form regressions like (4) is that the OLS coefficient estimates might be biased because the regressions do not adequately account for all relevant right-hand size variables (such as expectations about income growth; see also Appendix 2 for further discussion).

³²Interestingly, an auxiliary regression of income growth on the lagged saving rate in the spirit of Campbell (1987) yields statistically insignificant slope when post-1985 data are included (see Table 7 in Appendix 2).

of the financial sector continues improving, the saving rate is likely to remain high so long as uncertainty remains high and household wealth remains impaired (compared, at least, to its previous heights).

5 Structural Estimation

This section estimates the structural model of section 2 by minimizing the distance between the data on saving implied by the model and those observed in reality. The nonlinear least squares (NLLS) procedure we use has some advantages over the reduced-form regressions. Besides arguably being more immune to endogeneity and suitable for estimating structural parameters (such as the discount factor), it imposes on the data a structure that makes them easier to interpret. In particular, the model identifies a value for target wealth, which varies depending on the evolution of risk and credit conditions, and which can in principle be useful for identifying major deviations of actual wealth from the optimal level desired by consumers and gauging future trends in the saving rate. As Figure 2 documents, the structural model explicitly justifies and disciplines non-linearities, which can be important especially during turbulent times, when the shocks are large enough to move the system far from its steady state. In such times, estimation of linear or linearized models may be subject to substantial error.

5.1 Estimation Procedure

We assume households instantaneously observe exogenous movements in the three factors: wealth shocks m , unemployment risk \mathcal{U} and credit supply conditions CEA, and that they consider the shocks to \mathcal{U} and CEA to be permanent (and do not expect the shocks to wealth to be reversed).³³ Given these factors and the parameters, each period consumers re-optimize their consumption–saving choice (described in section 2). Collecting the parameters in the vector Θ and denoting the target wealth $\tilde{m}_t(\cdot)$ and the corresponding wealth gap $m_t - \tilde{m}_t$, the model implies a series of saving rates $s_t^{\text{theor}}(\Theta; m_t - \tilde{m}_t)$, which we match to those observed in the data, s_t^{meas} . Our estimates $\hat{\Theta}$ thus solve the following problem:

$$\hat{\Theta} = \arg \min \sum_{t=1}^T \left(s_t^{\text{meas}} - s_t^{\text{theor}} \left(\Theta; m_t - \tilde{m}(\bar{m}(\text{CEA}_t), \mathcal{U}(\mathbb{E}_t u_{t+4})) \right) \right)^2, \quad (5)$$

³³The assumption that households believe the shocks to be permanent is necessary for us to be able to use the tractable model we described earlier in the paper. While indefensible as a literal proposition (presumably nobody believes the unemployment rate will remain high forever), the high serial correlation of these variables means that the assumption may not be too objectionable. In any case, a model that incorporated more realistic descriptions of these processes would be much less transparent and might not be computationally feasible with present technology.

where the target wealth \bar{m} depends on the credit conditions and unemployment risk as described in section 2. In our baseline specification the parameter vector Θ consists of the discount factor β and the scaling constants for credit conditions and unemployment risk:

$$\Theta = \{\beta, \bar{\theta}_m, \theta_{\text{CEA}}, \bar{\theta}_U, \theta_u\}, \quad (6)$$

$$\bar{m}_t = \bar{\theta}_m + \theta_{\text{CEA}} \text{CEA}_t, \quad (7)$$

$$\bar{U}_t = \bar{\theta}_U + \theta_u \mathbb{E}_t u_{t+4}. \quad (8)$$

The re-scaling ensures that the unitless measure of credit conditions is re-normalized as a fraction of disposable income and that the expected unemployment rate is transformed into the model-compatible equivalent of *permanent* risk. The model implies that $\theta_{\text{CEA}} > 0$ and $\theta_u > 0$.

Minimization (5) is a non-linear least squares problem in which the standard asymptotic results apply. Standard errors for the estimated parameters are calculated using the delta method as follows.³⁴ Define the scores $q_t(\Theta) = (s_t^{\text{meas}} - s_t^{\text{theor}}(\Theta)) \frac{\partial s_t^{\text{theor}}(\Theta)}{\partial \Theta'}$ and the 5×5 matrices $E = \text{var}(q_t(\Theta))$ and $D = \mathbb{E} \frac{\partial q_t(\Theta)}{\partial \Theta'}$. The estimates have the asymptotic distribution:

$$T^{1/2}(\hat{\Theta} - \Theta) \rightarrow_d \mathcal{N}(0, D^{-1}ED^{-1}).$$

Because the saving function $s_t^{\text{theor}}(\Theta)$ is not available in the closed form, we calculate its partial derivatives numerically.

5.2 Results


Table 4 summarizes the calibration and the estimation results. The calibrated parameters—real interest rate $r = 0.04/4$, wage growth $\Delta W = 0.01/4$ and the coefficient of relative risk aversion $\rho = 2$ take their standard (quarterly) values and meet (together with the discount factor β) the conditions sufficient for the problem to be well-defined.

The discount factor $\beta = 1 - 0.0064 = 0.9936$, or 0.975 at annual frequency, lies in the standard range.

Figure 10 shows the estimated horizontal shift in the consumption function \bar{m}_t . The estimates of the scaling factors $\bar{\theta}_m$ and θ_{CEA} imply that \bar{m}_t varies roughly between 0 and $(-0.0071 + 5.2208)/4 = 1.3$, implying that financial deregulation resulted at its peak in an availability of credit in 2007 that was greater than credit availability at the beginning of our sample in 1966 by an amount equal to about 130% of annual income—not an unreasonable figure.

³⁴To construct the objective function (which we then minimize over Θ) we need to solve the consumer's optimization for each quarter. Because the calculation is computationally demanding, we cannot apply bootstrap to calculate standard errors. (The Shapiro–Wilk test does not reject normality of residuals.)


Figure 10 Extent of Credit Constraints \bar{m}_t (Fraction of Quarterly Disposable Income)



Notes: Shading—NBER recessions.

Sources: Federal Reserve, accumulated scores from the question on change in the banks' willingness to provide consumer installment loans from the Senior Loan Officer Opinion Survey on Bank Lending Practices, <http://www.federalreserve.gov/boarddocs/snloansurvey/>, authors' calculations.

Figure 11 Per Quarter Permanent Unemployment Risk \mathcal{U}_t



Notes: Shading—NBER recessions.


Sources: Thomson Reuters/University of Michigan Surveys of Consumers, <http://www.sca.isr.umich.edu/main.php>, Bureau of Labor Statistics, authors' calculations.

Figure 11 shows the estimated quarterly intensity of perceived permanent unemployment risk.

Figure 12 shows the fit of the structural model. In terms of \bar{R}^2 (Table 4), the model captures more than 80 percent of variation in the saving rate, doing only slightly worse than our baseline reduced-form model (whose \bar{R}^2 is roughly 0.9). The Mincer–Zarnowitz horse race between the models puts roughly 0.45 on the structural model (although the high standard error on the coefficient signals the high correlation between the two model-implied saving rates).

In principle, time variation in the fitted saving rate arises in our model exclusively due to the precautionary motive, which can be broken down into the three components: uncertainty, wealth, and credit conditions, as shown in Figure 13. Given the estimated parameters Θ (from Table 4) we sequentially switch off the uncertainty and


Figure 12 Fit of the Structural Model—Actual and Fitted PSR (Percent of Disposable Income)



Legend: Actual PSR: Thin black line, Structural model: Thick red/grey line. Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors' calculations.

Figure 13 Decomposition of Fitted PSR (Percent of Disposable Income)



Notes: Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors' calculations.

credit supply channels by setting the values of these series equal to their sample means. This means that, e.g., the difference between the model fitted series (red/grey line and the fitted series excluding uncertainty (black line) in Figure 13 is to be interpreted as the effects of *time variation* in unemployment risk \bar{U} (rather than total extent of saving caused by the existence of uncertainty).

While the wealth fluctuations do contribute to a good performance of the model at the business-cycle frequencies, the CEA is essential in capturing the trend decline in the PSR between the 1980s and the early 2000s. The principal role of cyclical fluctuations in uncertainty is to magnify the increases in the PSR during recessions, including the last one.

Table 5 replicates the estimates of Table 1 for the (artificial) saving rate series generated by the estimated structural model. The fact that the coefficient estimates closely mirror those obtained in actual time series documents that the structural

model captures well key features of the data on saving. Unsurprisingly, the standard errors are somewhat smaller than those in Table 1 and the \bar{R}^2 s are higher because the process of generating the artificial data by the model eliminates much of the noise (which is present in the actual data on the PSR).

6 Other Interpretations

There can be little doubt that factors aside from those included in our model have had some effect on U.S. saving dynamics over our sample period. For example, Sabelhaus and Song (2010) show a substantial decline in the size of both transitory and permanent shocks to income over the past 40 years; this should have led to a decline in precautionary saving that is probably not fully captured by the fact that our measure of unemployment risk is somewhat lower in the latter than in the earlier part of our sample. The shift from defined benefit to defined contribution pension plans may also have played an important role in changing the saving rate. And changing rates of taxation on capital income may also have played a role. Finally, the large increase in income inequality that has taken place since the mid-1970s might have been expected to result in an increase in the saving rate, since microeconomic evidence shows a robust tendency of higher-permanent-income households to save more (Carroll (2000); Dynan, Skinner, and Zeldes (2004)).

It seems unlikely, however, that such effects will be distinguishable from those are now included in our model using aggregate data of the kind employed in our paper, for the simple reason that the model's fit is already so good that the likelihood of achieving substantial further improvement by the addition of other variables seems remote. (For the version of the model that allows for the AR(1) component in saving implied by the theory, the R^2 is about 0.93). This does not, of course, mean that those other factors were not important; instead, it could be that the variables already included in our analysis are correlated with an omitted variable (say, the pension shift may be correlated with the availability of credit) to such an extent that aggregate data cannot be effectively used to discover the "real" contributions of the omitted versus the included variables. Real progress on such questions will likely require the use of good "natural experiments" in a microeconomic setting.

7 Conclusions

We find evidence that credit availability, shocks to household wealth, and movements in income uncertainty proxied by unemployment risk have all been important factors in driving U.S. household saving over the past 45 years. In particular, a relentless expansion of credit supply between the early-1980s and 2007 (likely largely reflecting


financial innovation and liberalization), along with higher asset values and consequent increases in net wealth (possibly also partly attributable to the credit boom) encouraged households to save less out of their disposable income. At the same time, the fluctuations in net wealth and labor income uncertainty, for instance during and after the burst of the information technology and credit bubbles of 2001 and 2007, can explain the bulk of business cycle fluctuations in personal saving.

We also find that other determinants of saving suggested by various literatures (e.g., fiscal deficits, demographics, income expectations) either work through the key factors above, are of second-order importance, or matter only during particular episodes. These findings are broadly in line with the complementary household-level evidence reported in Dynan and Kohn (2007), Moore and Palumbo (2010), Bricker, Bucks, Kennickell, Mach, and Moore (2011) and Petev, Pistaferri, and Eksten (2011).³⁵

Of course, all this evidence is based on historical data and, going forward, factors such as rapidly rising federal debt or the retirement of baby-boomers could yet lead to new structural shifts in household saving. But our results suggest that the personal saving rate in the pre-crisis period was artificially low because of the bubble in housing prices and the corresponding easy availability of credit. Neither of these factors seems likely to return soon, and since consensus forecasts suggest that the unemployment rate is likely to remain elevated for a long time, there seems to be little prospect that the personal saving rate will return to its low pre-crisis value anytime in the foreseeable future.

³⁵Dynan and Kohn (2007) find that data from the Federal Reserve's Survey of Consumer Finances (SCF) and the Michigan Survey of Consumer Sentiment show too little variation in the measures of impatience, risk aversion, expected income, interest rates and demographics to adequately explain the household indebtedness. In contrast, they argue that house prices and financial innovation have been important drivers of indebtedness. Moore and Palumbo (2010) document that the drop in consumer spending during the Great Recession was accompanied by significant erosions of home and corporate equity held by households. Using SCF data, Bricker, Bucks, Kennickell, Mach, and Moore (2011) document higher desired precautionary saving among most families during the Great Recession. Petev, Pistaferri, and Eksten (2011) discuss the following factors behind the observed changes in consumption during the Great Recession: the wealth effect, an increase in uncertainty and the credit crunch.

Figure 14 Alternative Measures of Credit Availability



Legend: Debt–disposable income ratio: Thin black line, CEA index: Thick red/grey line, the Abiad, Detragiache, and Tressel (2008) Index of Financial Liberalization: Dashed black line. Shading—NBER recessions.


Sources: Federal Reserve, accumulated scores from the question on change in the banks’ willingness to provide consumer installment loans from the Senior Loan Officer Opinion Survey on Bank Lending Practices, <http://www.federalreserve.gov/boarddocs/snloansurvey/>; Abiad, Detragiache, and Tressel (2008); Flow of Funds, Board of Governors of the Federal Reserve System.

Appendix 1: Comparison of Alternative Measures of Credit Availability

Figure 14 compares three measures of credit availability: our baseline CEA index, the Index of Financial Liberalization constructed of Abiad, Detragiache, and Tressel (2008) for a number of countries including the United States, and the ratio of household liabilities to disposable income.

The Abiad, Detragiache, and Tressel index is a mixture of indicators of financial development: credit controls and reserve requirements, aggregate credit ceilings, interest rate liberalization, banking sector entry, capital account transactions, development of securities markets and banking sector supervision. The correlation coefficient between this measure and CEA is about 90 percent.

Figure 15 Growth of Real Disposable Income (Percent)



Legend: BEA disposable income: Thick red/grey line, “Less cleaned” disposable income series: Thin black line, “More cleaned” disposable income series: Dashed black line. Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors’ calculations.


For comparison, the figure also includes the ratio of liabilities to disposable income (from the Flow of Funds), which is admittedly determined by the interaction between credit supply and demand.

Appendix 2: Stochastic Properties of Aggregate Disposable Income

Measurement of Disposable Income

This appendix investigates the properties of three measures of disposable income: the official series produced by the BEA and two alternative “cleaned” series, in which we try to exclude transitory income shocks due to temporary events, such as weather and fiscal policy.

Figure 16 Personal Saving Rate (Percent of Disposable Income)



Legend: BEA personal saving rate: Thick red/grey line, PSR calculated with the “less cleaned” income series: Thin black line, PSR calculated with the “more cleaned” income series: Dashed black line. Shading—NBER recessions.

Sources: Bureau of Economic Analysis, authors’ calculations.

Specifically, we have removed the following events from the official disposable income series using regressions:

- The dollar amounts of temporary rebate checks during 1975, 2008, and 2009 fiscal stimulus episodes.
- Dummies for the 20 costliest tropical cyclones using data from the National Weather Service.
- Dummies for quarters with unusually high or low cooling degree days, and unusually high or low heating degree days (the dummy has a value of 1 whenever the seasonally-adjusted series are more than 2 standard deviations above or below its mean).
- Dummies for quarters with unusually high or low national temperature, and unusually high or low precipitation (again, using the 2 standard deviations criterion).
- Separate dummies for snowstorms or heat waves which were deemed unusually extensive and damaging (these events do not necessarily overlap with the episodes identified from the national temperature and cooling/heating degree days data).

The “less cleaned” disposable income series removes from published data the contributions of stimulus and heating/cooling day extremes. The “more cleaned” series removes all the sources of transitory fluctuations outlined above.

Stochastic Properties of Disposable Income and Saving for a Rainy Day

The classic paper by Campbell (1987) has derived that the permanent income hypothesis implies that saving is negatively related to future expected income growth. This appendix investigates the univariate stochastic properties of disposable income and the relationship between saving and income, or the lack of it, in Tables 6 and 7, respectively.

Table 6 documents that all three disposable income series are statistically indistinguishable from a random walk. This means that the series are unpredictable using their own lags. In particular, for the income series in log-level, the first autocorrelations are very close to 1 and the augmented Dickey–Fuller test does not reject the null of a unit root. In contrast, for income growth, the first and other autocorrelations are zero, as also documented by the p values of the Box–Ljung Q statistic, and the ADF test (of course) strongly rejects a unit root.

Table 7 reports the estimates of α_1 the sensitivity of the saving rate to future income growth:

$$s_t = \alpha_0 + \alpha_1 \Delta y_{t+1} + \varepsilon_t, \tag{9}$$

which is motivated by Campbell (1987), who derives that under the permanent income hypothesis the coefficient α_1 is negative, as households save more when they are pessimistic about future income growth.

Overall, the estimates suggest that coefficient α_1 is statistically insignificant and small, especially when the full sample, 1966q2–2011q1, is used and when income growth Δy_{t+2}

Table 1 Preliminary Saving Regressions and the Time Trend
$$s_t = \gamma_0 + \gamma_m m_t + \gamma_{CEA} CEA_t + \gamma_{Eu} \mathbb{E}_t u_{t+4} + \gamma_t t + \gamma_{uC} (\mathbb{E}_t u_{t+4} \times CEA_t) + \varepsilon_t$$

Model	Time	Wealth	CEA	Un Risk	All 3	Baseline	Interact
γ_0	11.954*** (0.608)	25.202*** (1.727)	9.321*** (0.574)	8.241*** (0.420)	14.896*** (2.558)	15.226*** (2.157)	15.550*** (2.556)
γ_m		-2.606*** (0.319)			-1.124*** (0.423)	-1.183*** (0.347)	-1.368*** (0.456)
γ_{CEA}			-14.138*** (1.736)		-5.472*** (1.936)	-6.121*** (0.573)	-4.604*** (1.721)
γ_{Eu}				0.670*** (0.055)	0.316*** (0.117)	0.287*** (0.075)	0.385*** (0.108)
γ_t	-0.044*** (0.005)	-0.025*** (0.003)	0.042*** (0.011)	-0.048*** (0.002)	-0.005 (0.014)		0.004 (0.014)
γ_{uC}							-0.321** (0.158)
\bar{R}^2	0.703	0.846	0.825	0.881	0.895	0.895	0.899
F stat p val	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
DW stat	0.305	0.686	0.500	0.863	0.936	0.933	0.980

Notes: Estimation sample: 1966q2–2011q1. {*, **, ***} = Statistical significance at {10, 5, 1} percent. Newey–West standard errors, 4 lags.

enters the regression (9), which might be justified because of time aggregation issues. While there is some evidence of a negative coefficient in the pre-1985 sample (which overlaps with the sample 1953q2–1984q4 considered by Campbell (1987)), the relationship seems to break down in the past 20 years.

Table 2 Additional Saving Regressions

$$s_t = \gamma_0 + \gamma_m m_t + \gamma_{CEA} CEA_t + \gamma_{Eu} \mathbb{E}_t u_{t+4} + \gamma_\sigma \sigma_t + \gamma_s s_{t-1} + \gamma_d d_t + \dots \\ \dots + \gamma_r r_t + \gamma_{GS} GS_t + \gamma_{CS} CS_t + \varepsilon_t$$

Model	Baseline	Uncertainty	Lagged s_{t-1}	Debt	Full Controls	Post-1980	IV
γ_0	15.226*** (2.157)	15.080*** (2.180)	5.323*** (1.667)	13.884*** (2.100)	17.459*** (1.877)	16.692** (7.571)	21.323*** (2.746)
γ_m	-1.183*** (0.347)	-1.211*** (0.363)	-0.307 (0.222)	-0.803** (0.360)	-1.304*** (0.308)	-1.503 (1.248)	-2.022*** (0.492)
γ_{CEA}	-6.121*** (0.573)	-5.967*** (0.648)	-2.874*** (0.531)	-5.399*** (0.732)	-6.242*** (0.628)	-4.999** (2.000)	-5.846*** (1.166)
γ_{Eu}	0.287*** (0.075)	0.282*** (0.094)	0.143*** (0.053)	0.345*** (0.071)	0.117 (0.088)	0.298** (0.136)	0.084 (0.133)
γ_σ		0.257 (0.466)					
γ_s			0.574*** (0.072)				
γ_d				-1.905 (1.162)			
γ_r					0.129*** (0.043)		
γ_{GS}					-0.121 (0.081)		
γ_{CS}					-0.310** (0.138)		
$\gamma_{0post80}$						-1.479 (7.905)	
$\gamma_{mpost80}$						0.559 (1.289)	
$\gamma_{CEApost80}$						-2.350 (2.135)	
$\gamma_{Eupost80}$						-0.098 (0.162)	
\bar{R}^2	0.895	0.896	0.927	0.898	0.910	0.899	
F stat p val	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
F p val post 80						0.16665	
DW stat	0.933	0.940	2.134	0.924	0.954	0.967	
OID p val							0.740

Notes: Estimation sample: 1966q2–2011q1. {*, **, ***} = Statistical significance at {10, 5, 1} percent. Newey–West standard errors, 4 lags. CEA is the Credit Easing Accumulated Index, GS is the government saving as a fraction of GDP, CS is the corporate saving as a fraction of GDP. In model IV, m , CEA and $\mathbb{E}u$ are instrumented with lags 1 and 2 of m , $\mathbb{E}u$ and the Abiad, Detragiache, and Tressel (2008) Index of Financial Liberalization; the sample for the IV model is 1973q1–2005q4 (as in Abiad, Detragiache, and Tressel (2008)). OID p val denotes the p-value from the Hansen’s J statistic for overidentification.

Table 3 Personal saving rate—Actual and explained change, 2007–2010

Variable	Baseline	Interact	Actual Δs_t
$\gamma_m \times \Delta m_t$	$-1.18 \times -1.39 = 1.64$	$-1.37 \times -1.39 = 1.90$	
$\gamma_{CEA} \times \Delta CEA_t$	$-6.12 \times -0.11 = 0.64$	$-4.60 \times -0.11 = 0.48$	
$\gamma_{Eu} \times \Delta \mathbb{E}_t u_{t+4}$	$0.29 \times 4.33 = 1.24$	$0.38 \times 4.33 = 1.67$	
$\gamma_{uC} \times \Delta(\mathbb{E}_t u_{t+4} \times CEA_t)$		$-0.32 \times 3.33 = -1.07$	
Explained Δs_t	3.53	2.98	2.93

Table 4 Calibration and Structural Estimates

$$\begin{aligned}
s_t^{\text{theor}} &= s_t^{\text{theor}}(\Theta; m_t - \check{m}(\bar{m}_t, \mathcal{U}_t)), \\
\bar{m}_t &= \bar{\theta}_m + \theta_{\text{CEA}} \text{CEA}_t, \\
\mathcal{U}_t &= \bar{\theta}_{\mathcal{U}} + \theta_u \mathbb{E}_t u_{t+4}.
\end{aligned}$$

Parameter	Description	Value
Calibrated Parameters		
r	Interest Rate	0.04/4
ΔW	Wage Growth	0.01/4
ρ	Relative Risk Aversion	2
Estimated Parameters $\Theta = \{\beta, \bar{\theta}_m, \theta_{\text{CEA}}, \bar{\theta}_{\mathcal{U}}, \theta_u\}$		
β	Discount Rate	$1 - 0.0064^{***}$ (0.0018)
$\bar{\theta}_m$	Scaling of \bar{m}_t	0.0072 (0.0206)
θ_{CEA}	Scaling of \bar{m}_t	5.2215^{***} (0.1396)
$\bar{\theta}_{\mathcal{U}}$	Scaling of \mathcal{U}_t	5.3758×10^{-5} (8.4334×10^{-5})
θ_u	Scaling of \mathcal{U}_t	0.0363 (0.1227)
\bar{R}^2		0.821
DW stat		0.950

Notes: Quarterly calibration. Estimation sample: 1966q2–2011q1. {*, **, ***} = Statistical significance at {10, 5, 1} percent. Standard errors (in parentheses) were calculated with the delta method.

Table 5 Preliminary Saving Regressions and the Time Trend—Saving Rate
Generated by the Structural Model

$$s_t = \gamma_0 + \gamma_m m_t + \gamma_{CEA} CEA_t + \gamma_{Eu} \mathbb{E}_t u_{t+4} + \gamma_t t + \gamma_{uC} (\mathbb{E}_t u_{t+4} \times CEA_t) + \varepsilon_t$$

Model	Time	Wealth	CEA	Un Risk	All 3	Baseline	Interact
γ_0	11.955*** (0.502)	23.765*** (1.355)	9.354*** (0.410)	8.422*** (0.160)	13.031*** (0.720)	13.357*** (0.634)	13.423*** (0.656)
γ_m		-2.327*** (0.251)			-0.790*** (0.120)	-0.848*** (0.105)	-0.936*** (0.108)
γ_{CEA}			-13.821*** (1.124)		-5.846*** (0.594)	-6.486*** (0.141)	-5.327*** (0.467)
γ_{Eu}				0.633*** (0.024)	0.328*** (0.035)	0.299*** (0.019)	0.369*** (0.030)
γ_t	-0.044*** (0.004)	-0.027*** (0.002)	0.040*** (0.007)	-0.048*** (0.001)	-0.005 (0.004)		0.000 (0.003)
γ_{uC}							-0.192*** (0.037)
\bar{R}^2	0.799	0.929	0.931	0.979	0.993	0.992	0.994
F stat p val	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
DW stat	0.053	0.220	0.095	0.387	0.721	0.714	0.994

Notes: Estimation sample: 1966q2–2011q1. {*, **, ***} = Statistical significance at {10, 5, 1} percent. Newey–West standard errors, 4 lags.

Table 6 Univariate Properties of Disposable Income and Personal Saving Rate

Series	Official BEA	Less Cleaned	More Cleaned
	Disposable Income—Log-level		
First Autocorrelation	0.983	0.983	0.983
Box–Ljung Q stat, p value	0.000	0.000	0.000
Augmented Dickey–Fuller test, p value	0.505	0.515	0.501
	Disposable Income—Growth Rate		
First Autocorrelation	−0.043	−0.033	−0.024
Box–Ljung Q stat, p value	0.604	0.446	0.334
Augmented Dickey–Fuller test, p value	0.000	0.000	0.000
	Personal Saving Rate		
First Autocorrelation	0.953	0.953	0.952
Box–Ljung Q stat, p value	0.000	0.000	0.000
Augmented Dickey–Fuller test, p value	0.628	0.600	0.539

Notes: Box–Ljung statistics: 8 lags, ADF test: 4 lags.

Table 7 Campbell (1987) Saving for a Rainy Day Regressions

Series	Official BEA	Less Cleaned	More Cleaned
Full Sample: 1966q2–2011q1			
$s_t = \alpha_0 + \alpha_1 \Delta y_{t+1} + \varepsilon_t$			
α_1	−0.046 (0.052)	−0.054 (0.050)	−0.065 (0.051)
\bar{R}^2	−0.002	−0.000	0.002
$s_t = \alpha_0 + \alpha_1 \Delta y_{t+2} + \varepsilon_t$			
α_1	0.017 (0.050)	0.009 (0.047)	−0.009 (0.047)
\bar{R}^2	−0.005	−0.006	−0.006
Pre-1985 Sample: 1966q2–1984q4			
$s_t = \alpha_0 + \alpha_1 \Delta y_{t+1} + \varepsilon_t$			
α_1	−0.108*** (0.031)	−0.105*** (0.031)	−0.117*** (0.033)
\bar{R}^2	0.143	0.128	0.150
$s_t = \alpha_0 + \alpha_1 \Delta y_{t+2} + \varepsilon_t$			
α_1	−0.056 (0.039)	−0.060* (0.036)	−0.083** (0.033)
\bar{R}^2	0.029	0.034	0.070

Notes: {*, **, ***} = Statistical significance at {10, 5, 1} percent. Newey–West standard errors, 4 lags.

References

- ABIAD, ABDUL, ENRICA DETRAGIACHE, AND THIERRY TRESSEL (2008): “A New Database of Financial Reforms,” working paper 266, International Monetary Fund.
- AIYAGARI, S. RAO (1993): “Optimal Capital Income Taxation with Incomplete Markets and Borrowing Constraints,” Federal Reserve Bank of Minneapolis Working Paper Number 508.
- ARON, JANINE, JOHN V. DUCA, JOHN MUELLBAUER, KEIKO MURATA, AND ANTHONY MURPHY (2011): “Credit, Housing Collateral, and Consumption: Evidence from Japan, the U.K., and the U.S.,” *Review of Income and Wealth*.
- BAKER, SCOTT R., NICHOLAS BLOOM, AND STEVEN J. DAVIS (2011): “Measuring Economic Policy Uncertainty,” mimeo, Stanford University and University of Chicago.
- BEWLEY, TRUMAN (1977): “The Permanent Income Hypothesis: A Theoretical Formulation,” *Journal of Economic Theory*, 16, 252–292.
- BLINDER, ALAN S., AND ANGUS S. DEATON (1985): “The Time Series Consumption Function Revisited,” *Brookings Papers on Economic Activity*, 1985:2, pp. 465–521.
- BLOOM, DAVID E., DAVID CANNING, RICHARD K. MANSFIELD, AND MICHAEL MOORE (2007): “Demographic Change, Social Security Systems and Savings,” *Journal of Monetary Economics*, 54, 92–114.
- BLOOM, NICHOLAS, MAX FLOETOTTO, AND NIR JAIMOVICH (2009): “Really Uncertain Business Cycles,” mimeo, Stanford University.
- BOSWORTH, BARRY, AND GABRIEL CHODOROW-REICH (2007): “Saving and Demographic Change: The Global Dimension,” working paper 2, Center for Retirement Research, Boston College.
- BRICKER, JESSE, BRIAN BUCKS, ARTHUR KENNICKELL, TRACI MACH, AND KEVIN MOORE (2011): “Surveying the Aftermath of the Storm: Changes in Family Finances from 2007 to 2009,” FEDS working paper 17, Federal Reserve Board.
- CAMPBELL, JEFFREY R., AND ZVI HERCOWITZ (2009): “Welfare Implications of the Transition to High Household Debt,” *Journal of Monetary Economics*, 56(1), 1–16.
- CAMPBELL, JOHN Y. (1987): “Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis,” *Econometrica*, 55, 1249–1273.
- CARROLL, CHRISTOPHER D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 1992(2), 61–156, <http://econ.jhu.edu/people/ccarroll/BufferStockBPEA.pdf>.

- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), 1–55.
- (2000): “Why Do the Rich Save So Much?,” in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. by Joel B. Slemrod. Harvard University Press, <http://econ.jhu.edu/people/ccarroll/Why.pdf>.
- (2001): “A Theory of the Consumption Function, With and Without Liquidity Constraints,” *Journal of Economic Perspectives*, 15(3), 23–46, <http://econ.jhu.edu/people/ccarroll/ATheoryv3JEP.pdf>.
- CARROLL, CHRISTOPHER D., AND OLIVIER JEANNE (2009): “A Tractable Model of Precautionary Reserves, Net Foreign Assets, or Sovereign Wealth Funds,” *NBER Working Paper Number 15228*, <http://econ.jhu.edu/people/ccarroll/papers/cjsOE>.
- CARROLL, CHRISTOPHER D., AND MILES S. KIMBALL (1996): “On the Concavity of the Consumption Function,” *Econometrica*, 64(4), 981–992, <http://econ.jhu.edu/people/ccarroll/concavity.pdf>.
- CARROLL, CHRISTOPHER D., MISUZU OTSUKA, AND JIRI SLACALEK (2011): “How Large Are Financial and Housing Wealth Effects? A New Approach,” *Journal of Money, Credit, and Banking*, 43(1), 55–79.
- CARROLL, CHRISTOPHER D., AND PATRICK TOCHE (2009): “A Tractable Model of Buffer Stock Saving,” *NBER Working Paper Number 15265*, <http://econ.jhu.edu/people/ccarroll/papers/ctDiscrete>.
- CASE, KARL E., JOHN M. QUIGLEY, AND ROBERT J. SHILLER (2005): “Comparing Wealth Effects: The Stock Market Versus the Housing Market,” *Advances in Macroeconomics*, 5(1), 1–32.
- CHALLE, EDOUARD, AND XAVIER RAGOT (2012): “Precautionary Saving over the Business Cycle,” mimeo, Paris School of Economics.
- CHRISTIANO, LAURENCE J., MARTIN EICHENBAUM, AND CHARLES L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1), 1–45.
- DAVIS, MORRIS, AND MICHAEL PALUMBO (2001): “A Primer on the Economics and Time Series Econometrics of Wealth Effects,” *Federal Reserve Board Finance and Economics Discussion Papers 2001-09*.
- DEATON, ANGUS (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–48.
- 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.

- DEUTSCHE BANK SECURITIES (2012): “Global Economic Perspectives: Risks to US Consumer Spending Overrated,” *Deutsche Bank Global Markets Research: Macro*.
- DUCA, JOHN V., JOHN MUELLBAUER, AND ANTHONY MURPHY (2010): “Credit Market Architecture and the Boom and Bust in the U.S. Consumption,” mimeo, University of Oxford.
- (2011): “House Prices and Credit Constraints: Making Sense of the US Experience,” *The Economic Journal*, 121(552), 533–551.
- DYNAN, KAREN E., DOUGLAS W. ELMENDORF, AND DANIEL E. SICHEL (2006): “Can Financial Innovation Help to Explain the Reduced Volatility of Economic Activity?,” *Journal of Monetary Economics*, 53, 123–150.
- DYNAN, KAREN E., AND DONALD L. KOHN (2007): “The Rise in US Household Indebtedness: Causes and Consequences,” in *The Structure and Resilience of the Financial System*, ed. by Christopher Kent, and Jeremy Lawson, pp. 84–113. Reserve Bank of Australia.
- DYNAN, KAREN E., JONATHAN SKINNER, AND STEPHEN P. ZELDES (2004): “Do the Rich Save More?,” *Journal of Political Economy*, 112(2), 397–444.
- EGGERTSSON, GAUTI B., AND PAUL KRUGMAN (2011): “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *Manuscript, NBER Summer Institute*.
- FERNANDEZ-CORUGEDO, EMILIO, AND JOHN MUELLBAUER (2006): “Consumer Credit Conditions in the United Kingdom,” working paper 314, Bank of England.
- FISHER, IRVING (1930): *The Theory of Interest*. MacMillan, New York.
- FRIEDMAN, MILTON A. (1957): *A Theory of the Consumption Function*. Princeton University Press.
- GREEN, RICHARD K., AND SUSAN M. WACHTER (2007): “The Housing Finance Revolution,” in *Housing, Housing Finance and Monetary Policy*, pp. 21–67. Jackson Hole Symposium, Federal Reserve Bank of Kansas City.
- GUERRIERI, VERONICA, AND GUIDO LORENZONI (2011): “Credit Crises, Precautionary Savings and the Liquidity Trap,” *Manuscript, MIT Department of Economics*.
- HALL, ROBERT E. (2011): “The Long Slump,” AEA Presidential Address, ASSA Meetings, Denver.
- KRUGMAN, PAUL (2012): “Postmodern Business Cycles,” *New York Times*, <http://krugman.blogs.nytimes.com/2012/01/27/postmodern-business-cycles/>.

- LETTAU, MARTIN, AND SYDNEY LUDVIGSON (2004): “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption,” *American Economic Review*, 94(1), 276–299, <http://www.jstor.org/stable/3592779>.
- LUDVIGSON, SYDNEY, AND CHARLES STEINDEL (1999): “How Important Is the Stock Market Effect on Consumption?,” *Economic Policy Review*, 5(2), 29–51.
- MCCARTHY, JONATHAN, AND RICHARD W. PEACH (2002): “Monetary Policy Transmission to Residential Investment,” FRBNY Economic Policy Review 1, Federal Reserve Bank of New York.
- MERTON, ROBERT C. (1969): “Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case,” *Review of Economics and Statistics*, 50, 247–257.
- MODY, ASHOKA, FRANZISKA OHNSORGE, AND DAMIANO SANDRI (2012): “Precautionary Saving in the Great Recession,” Discussion paper.
- MOORE, KEVIN B., AND MICHAEL G. PALUMBO (2010): “The Finances of American Households in the Past Three Recessions: Evidence from the Survey of Consumer Finances,” Finance and Economics Discussion Series 06, Federal Reserve Board.
- MUELLBAUER, JOHN N. (2007): “Housing, Credit and Consumer Expenditure,” in *Housing, Housing Finance and Monetary Policy*, pp. 267–334. Jackson Hole Symposium, Federal Reserve Bank of Kansas City.
- NAKAMURA, LEONARD I., AND THOMAS STARK (2007): “Mismeasured Personal Saving and the Permanent Income Hypothesis,” *Working Papers*, <http://philadelphiafed.org/research-and-data/publications/working-papers/2007/wp07-8.pdf>.
- PARKER, JONATHAN A. (2000): “Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate,” in *NBER Macroeconomics Annual 1999*, ed. by Ben S. Bernanke, and Julio J. Rotemberg, vol. 14, pp. 317–387. NBER.
- PETEV, IVAYLO, LUIGI PISTAFERRI, AND ITAY SAPORTA EKSTEN (2011): “Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects,” in *The Great Recession*, ed. by David B. Grusky, Bruce Western, and Christopher Wimer, pp. 161–196. Russell Sage Foundation.
- SABELHAUS, JOHN, AND JAE SONG (2010): “The Great Moderation in Micro Labor Earnings,” *Journal of Monetary Economics*, 57(4), 391–403, <http://ideas.repec.org/a/eee/moneco/v57y2010i4p391-403.html>.
- SAMUELSON, PAUL A. (1969): “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” *Review of Economics and Statistics*, 51, 239–46.
- SKINNER, JONATHAN (1988): “Risky Income, Life-Cycle Consumption, and Precautionary Saving,” *Journal of Monetary Economics*, 22, 237–55.

——— (1996): “Is Housing Wealth a Sideshow?,” in *Advances in the Economics of Aging*, ed. by David A. Wise, pp. 241–272. University of Chicago Press.

ZELDES, STEPHEN P. (1989): “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *Quarterly Journal of Economics*, 104(2), 275–98.