

Carroll, Christopher D.; Otsuka, Misuzu; Slacalek, Jirka

Working Paper

How large is the housing wealth effect? A new approach

Working Paper, No. 535

Provided in Cooperation with:

Department of Economics, The Johns Hopkins University

Suggested Citation: Carroll, Christopher D.; Otsuka, Misuzu; Slacalek, Jirka (2006) : How large is the housing wealth effect? A new approach, Working Paper, No. 535, The Johns Hopkins University, Department of Economics, Baltimore, MD

This Version is available at:

<https://hdl.handle.net/10419/72036>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

How Large Are Housing and Financial Wealth Effects? A New Approach

July 4, 2010

Christopher D. Carroll¹
JHU

Misuzu Otsuka²
OECD

Jiri Slacalek³
ECB

Abstract

This paper presents a simple new method for measuring ‘wealth effects’ on aggregate consumption. The method exploits the stickiness of consumption growth (sometimes interpreted as reflecting consumption ‘habits’) to distinguish between immediate and eventual wealth effects. In U.S. data, we estimate that the immediate (next-quarter) marginal propensity to consume from a \$1 change in housing wealth is about 2 cents, with a final eventual effect around 9 cents, substantially larger than the effect of shocks to financial wealth. We argue that our method is preferable to cointegration-based approaches, because neither theory nor evidence supports faith in the existence of a stable cointegrating vector.

Keywords Housing Wealth, Wealth Effect, Consumption Dynamics, Asset Prices

JEL codes E21, E32, C22

Published in Journal of Money, Credit, and Banking, Vol. 43, No. 1 (February 2011), pp 55–79

PDF: <http://econ.jhu.edu/people/ccarroll/papers/cosWealthEffects.pdf>

Web: <http://econ.jhu.edu/people/ccarroll/papers/cosWealthEffects/>

Archive: <http://econ.jhu.edu/people/ccarroll/papers/cosWealthEffects.zip>
(Contains data and estimation software producing paper’s results)

¹Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD, <http://econ.jhu.edu/people/ccarroll/>, ccarroll@jhu.edu ²Otsuka: Organisation for Economic Co-operation and Development, Paris, France, misuzu.otsuka@oecd.org ³Slacalek: European Central Bank, Frankfurt am Main, Germany, <http://www.slacalek.com/>, jiri.slacalek@ecb.europa.eu.

This paper is based on material originally prepared for an Academic Consultants’ meeting of the Board of Governors of the Federal Reserve System, January 30, 2004; for the version of the paper originally presented at the Fed, see Carroll (2004). We thank the discussant of that paper, Robert Shiller, as well as Pok-sang Lam (the

1 Introduction

Conventional wisdom says that the response of household spending to a shock to wealth (the ‘wealth effect’) has historically been around 3 to 5 cents on the dollar in the U.S.¹ However, much of the evidence for this proposition comes from ‘cointegrating’ models that regress the level of consumption on the levels of wealth and income.^{2,3}

We argue that cointegration methods are problematic for estimating wealth effects, for at least two reasons.⁴ First, basic consumption theory does not imply the existence of a stable cointegrating vector; in particular, a change in the long-run growth rate or the long-run interest rate should change the relationship between consumption, income, and wealth.⁵ Second, even if changes to the cointegrating vector are ruled out by assumption, changes in any other feature of the economy relevant for the consumption/saving decision can generate such long-lasting dynamics that hundreds or thousands of years of data should be required to obtain reliable estimates of that vector.

Even for the U.S., the technological leader and therefore the most stable advanced country in the modern era, the 50 year span of available data has seen major changes in productivity growth, interest tax rates, demographics, financial markets, social insurance, and every other aspect of reality that theory says should matter for consumption (not to mention fundamental changes in measurement methods for the underlying NIPA data).

¹This statement applies to the macroeconomics literature, cf. Davis and Palumbo (2001) and references therein. Results using microeconomic data are more heterogeneous; see Khalifa (2004) for capsule summaries of the literature at the time of the original drafting of this paper, and section 3.3 below for further discussion.

²This method was popular in the first generation of Keynesian econometric models from the 1950s; in the 1970s it was given rigorous econometric foundations and dubbed the ‘cointegrating’ approach, or (by a different tribe) the ‘error correction’ approach. It gained further respectability when Campbell and Mankiw (1989) developed a theoretical model that assumed that the ratio of human wealth to nonhuman wealth was stationary (although the theory they used provides no reason why that ratio should be stable; stability was an auxiliary assumption rather than an implication of their model).

³The number of papers applying the cointegration methodology has recently risen considerably. References include Bertaut (2002), Byrne and Davis (2003), Fernandez-Corugedo, Price, and Blake (2003), Pichette and Tremblay (2003), Catte, Girouard, Price, and Andre (2004), Lettau and Ludvigson (2004), Hamburg, Hoffmann, and Keller (2005).

⁴Any purported estimation of ‘wealth effects’ should reflect an acute awareness of the possibility (maybe, likelihood) that the putative ‘effect’ actually captures the influence of some omitted variable, like credit conditions (cf. Muellbauer (2007), Aron, Duca, Muellbauer, Murata, and Murphy (2008)) or growth expectations (Disney, Gathergood, and Henley (2008)). This footnote performs that duty; we mostly refrain from further apologies in the remainder of the text. See Calomiris, Longhofer, and Miles (2009) for such a critique.

⁵See Rudd and Whelan (2006) and Slacalek (2009) for empirical evidence of instability in the cointegrating vector.

editor), two anonymous referees, Ben Bernanke, Alan Greenspan, Robert Hall, Nicholas Souleles, David Wyss, Stephen Zeldes, the staff of the Fed, and other attendees at the meeting for very valuable feedback. The data and econometric programs that generated all of the results in this paper are available in the archive (url above). We are also grateful to Emmanuel de Veirman and Sherif Khalifa for excellent research assistance. The views presented in this paper are those of the authors, and should not be attributed to the Organisation for Economic Co-operation and Development or the European Central Bank.

Motivated by these concerns, we introduce an alternative methodology for estimating wealth effects. The method’s foundation derives from the recent literature documenting substantial ‘excess smoothness’ (or ‘stickiness’) in consumption growth, relative to the benchmark random walk model.⁶ Our model can be thought of as a proposal for unifying the ‘stickiness’ and the ‘wealth effects’ literatures, by resolving wealth effects into two key aspects: Speed and strength.

Our measure of ‘speed’ is meant to distill and quantify the core point of the stickiness literature: Consumption responds to shocks more slowly than implied by the random-walk benchmark. Given an estimated ‘speed,’ our measure of the ‘strength’ of wealth effects is thus dependent on the horizon; our ‘speed’ estimates imply that the immediate spending effects of wealth fluctuations are much smaller than the eventual effects.⁷ In particular, we find that the immediate (next-quarter) marginal propensity to consume (MPC) from a \$1 change in housing wealth is about 2 cents, with an eventual effect amounting to 9 cents. Consistent with several other recent studies, we find a housing wealth effect that is larger than the financial wealth effect, which we estimate to be about 6 cents on the dollar.

These differing estimates suggest that markets and policymakers worried about wealth effects may need to pay careful attention not only to the size of overall changes in total wealth but also to how those wealth changes break down between different asset classes.

2 A Theoretical Sketch

This section uses a simple model of consumption to illustrate why a cointegrating approach may not correctly identify the wealth effect (even when it exists), and shows how the new method that we propose addresses this shortcoming.

2.1 The Frictionless Model

In the benchmark perfect foresight model with no uncertainty, perfect capital markets, homogenous consumers and no bequest motive, steady-state consumption is proportional to overall resources. If total resources are the sum of nonmarket (human) wealth \mathbf{H} and market wealth \mathbf{B} , spending is given by

$$\mathbf{C}_t = (\mathbf{H}_t + \mathbf{B}_t)\kappa, \tag{1}$$

⁶The empirical literature on the sluggishness of consumption starts with Flavin (1981) and Campbell and Deaton (1989). Recent research suggests that this stylized fact can be alternatively explained by habit formation (see, e.g., Fuhrer, 2000; Chetty and Szeidl, 2005; Sommer, 2007) or inattentiveness (see, e.g., Carroll and Slacalek, 2006; Reis, 2006 and Sims, 2003).

⁷By ‘eventual’ we mean the effects after a few years; we are skeptical that our methods (or anyone else’s) can extract much true ‘long run’ (infinite horizon) information from macroeconomic data, for all the reasons mentioned above.

where κ is a constant determined by preferences and the after-tax interest factor $R = 1 + r$ (assumed here to be constant).⁸ In the version of the model with infinitely lived consumers, constant relative risk aversion ρ , and the discount factor $\beta = \frac{1}{1+\vartheta}$ the MPC takes the explicit form

$$\kappa = (R - (R\beta)^{1/\rho})/R.$$

If labor income is expected to grow at a constant rate g , then (in the continuous-time approximation) H_t will be given by $P_t/(r - g)$ where P_t is the current value of ‘permanent’ labor income, and (1) becomes

$$C_t = \left(\frac{\kappa}{r - g} \right) P_t + \kappa B_t. \quad (2)$$

This model can be extended to the case with i.i.d. interest rates (cf. Merton (1969) and Samuelson (1969)), which adds a term to the formula for κ but does not change the structure of (2). In this case the stochastic interest rate would result in ‘wealth shocks’ to B_t .

The cointegrating approach to estimating the model would be to assume that consumption is determined by an equation like (2) plus an error perhaps reflecting transitory shocks to consumption or measurement problems, leading to a regression of the form

$$C_t = \eta_0 + \eta_1 Y_t + \eta_2 B_t + \varepsilon_t$$

under the assumption that current income Y_t proxies for P_t and with the hope that the coefficient η_2 will uncover the MPC out of wealth, κ . (This is, of course, a simplification of actual practice, but it captures the essence of the method.)

Unfortunately, almost any attempt to make the model more realistic destroys the prediction that there is a time-invariant ‘wealth effect’ coefficient κ . In the simple model sketched above, it is clear that any sustained change in g , r , ρ , or ϑ during the estimation interval would pose serious problems because no time-invariant κ exists even under the usual maintained assumption that movements in B_t represent exogenous shocks; this point holds with even greater force if those ‘wealth shocks’ to B are correlated with persistent movements in g , r , ρ , or ϑ (as asset pricing theory suggests they will be).⁹

Even if we assume perpetual constancy in g , r , ρ , and ϑ , the model’s prediction

⁸The derivations in this section are standard, and can be found, for example, in the lecture notes on the first author’s web page.

⁹It is plausible, however, to suppose that these parameters are likely to change only gradually over time. Thus, any change in the cointegrating vector is likely to be gradual. Under these circumstances, an estimation method that allows for the cointegrating vector to evolve gradually might actually perform reasonably well; to see an attempt at an estimation along these lines, see the SMART estimation method explored in Aron, Duca, Muellbauer, Murata, and Murphy (2008). It would be interesting to test how well this method performs under plausible assumptions about the degree of drift in the coefficients on the cointegrating vector, but such an exercise is beyond the scope of the present paper.

of a time-invariant κ can be destroyed by the introduction of labor income uncertainty, time-varying after-tax interest rates, demographics, or many other real-world complications.

Such concerns are given further force by the large econometric literature on ‘spurious significance’ that can result from regressing non-stationary variables on each other, the upshot of which is that econometric tests may appear to detect a significant relationship even when the variables are actually independent.

As a specific illustration of the problem we are concerned about, consider the following scenario, illustrated in figures 1(a) and 1(b). (Plain solid lines show the reactions of consumption and wealth (normalized by income; hence nonbold) for the ‘frictionless’ model sketched above.) The economy starts in period 1 in a steady state balanced growth equilibrium in which $B = 0$ and $C = 1$, and stays in that equilibrium for 4 periods. In period 5 it is hit with a 1-unit positive shock to wealth so that $B_5 = 1$. The economy evolves with no further shocks for $n = 20$ quarters, then in period 25 experiences a permanent increase in income growth g .¹⁰ The simulation runs for another 20 periods, ending in period 45.¹¹

Consumption adjusts upward immediately (in period 5) to the period-5 wealth shock; both consumption and wealth remain constant thereafter. Thus, the size of the ‘wealth effect’ κ on consumption can be measured directly by comparing the change in consumption to the change in wealth. This is the ‘best-case scenario’ for measuring a wealth effect.

The expected growth rate of income g is the object that we modify in order to produce our experiment’s second shock. When this positive growth shock hits, consumption jumps up to a much higher level; but nonhuman wealth B begins to fall, because, now, spending exceeds income. This dissaving reflects the ‘human wealth’ effect emphasized by Summers (1981): When consumers become more optimistic about future income growth, they start spending today on the basis of their anticipated future riches. But as consumption continues to exceed current income, the ratio of nonhuman wealth to income B declines, and the ratio of consumption to income C also declines. Both C and B thus embark on downward trajectories after the shock; but B falls starting from its pre-shock level, while C starts falling only after having made a one-time upward leap because of the human wealth effect.

This experiment provides a clear example in which, even though a ‘marginal propensity to consume out of wealth’ unambiguously exists ($\kappa \approx 0.014$) in the

¹⁰The size of the first shock (the wealth shock) is arbitrarily set equal to 1 unit of permanent income. The income growth shock is an increase from 1.5 to 2.5 percent per year, roughly matching the rise in the ten-year productivity growth rate expected by the respondents to the *Survey of Professional Forecasters* between the early 1990s and the year 2001. (The annualized ten-year ahead forecast of productivity growth increased from 1.52 percent in 1996 to 2.48 percent in 2001, see <http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/data-files/PROD10/>.)

¹¹These simulations are described in more detail in the document `simulationsWithStickyC.pdf` in the replication archive.

underlying structural model, it cannot be uncovered by estimating a supposed ‘cointegrating’ regression of consumption on wealth. Indeed, the coefficient obtained from such a regression would actually be negative, because after the positive income-growth shock, consumption is higher on average than before the shock, while average wealth after the shock is lower than before the shock.

Table 1 presents the quantitative results for the ($n = 20$) experiment illustrated in the figures, as well as for similar experiments with 40 and 60 quarters.¹² As long as the shocks occur more frequently than about every 15 years (= 60 quarters), regressions of consumption C on wealth B estimate a negative wealth effect, even though the true parameter of interest is about 0.014. While cointegrating regressions eventually do provide consistent estimates as n approaches infinity and if there are no more shocks, Table 1 suggests that convergence of the cointegrating estimates to the truth is likely to be too slow to make cointegration estimation a reliable method of uncovering structural parameters (even if they exist) over the (relatively) short spans of time captured in actual empirical macroeconomic datasets.

2.2 The Sticky Expectations Model

This section argues that if there is a reliable degree of ‘stickiness’ in consumption growth, an estimation method that relies upon that stickiness to estimate wealth effects using high- and medium-frequency data is less likely to be led astray by a ‘regime change’ (like the one examined above) than a full-sample estimation technique like cointegration estimation.

Consumption habits are the leading explanation for sluggishness in aggregate consumption. But an alternative explanation is that households may be mildly *inattentive* to macroeconomic developments—for example, some households may not immediately notice shocks to aggregate macroeconomic indicators such as productivity growth or the unemployment rate. Carroll and Slacalek (2006) simulate an economy consisting of a continuum of such inattentive but otherwise-standard consumers with Constant Relative Risk Aversion utility, each of whom updates the information about his permanent income with probability Π in each period. They show that the change in the log of aggregate consumption, $\Delta \log \mathbf{C}_t$, approximately follows an autoregressive AR(1) process, whose autocorrelation coefficient approximates the share of consumers $(1 - \Pi)$ who do not have up-to-date information:

$$\Delta \log \mathbf{C}_t = \mu + \underbrace{(1 - \Pi)}_{\equiv \chi} \Delta \log \mathbf{C}_{t-1} + \varepsilon_t. \quad (3)$$

Exactly the same approximation can be obtained for some models of habit-forming consumers, e.g., Muellbauer (1988) and Dynan (2000), but the coefficient χ in those

¹²The table presents estimates for the frictionless model; similar estimates can be found for model with inattentive consumers described below.

models measures the intensity of the habit motive. When we estimate a model of the form of equation (3), our estimates cannot distinguish between these alternative hypotheses about the reason for stickiness.¹³ But, from the standpoint of forecasting aggregate consumption dynamics, it may not matter whether the right explanation of stickiness is habits or inattention.

As an illustration of how our estimation method works, consider the behavior of aggregate consumption and wealth in the same economy described in the previous section, except that consumption is now that of inattentive households who update their information on average once a year (i.e., $\Pi = 0.25$ or $\chi = 0.75$). The dashed lines in Figures 1(a) and 1(b) show the gradual adjustment of the two variables, which occurs because some consumers remain unaware of the shocks for several quarters. This sluggishness provides an informative signal to identify parameters of interest.

Table 2 uses an equation like the one we will estimate empirically below,

$$\Delta \mathbf{C}_t = \chi \mathbb{E}_{t-2} \Delta \mathbf{C}_{t-1} + \alpha \Delta \mathbf{B}_{t-1} + \varepsilon_t,$$

to estimate the strength of the impact of wealth α and the speed χ in simulated data.¹⁴ (We explain below why we use $\Delta \mathbf{B}_{t-1}$ rather than $\Delta \mathbf{B}_t$ as the second regressor.) The regression captures the essence of our estimation approach, whose empirical implementation is detailed in section 3. The two key findings in the table are: (i) the stickiness parameter χ lies close to its true value, and (ii) the estimates of the wealth effect are broadly in line with the ‘true’ value calculated from the calibrated parameters, 0.014. As for the latter, the estimated wealth effect decreases somewhat for $n = 60$, which suggests that enough shocks are needed to identify the parameters; but, shocks presumably arrive in the real world more often than once every 60 periods (15 years), so this finding is not especially troubling for our hopes of estimating the model with empirical data.

We should emphasize here that the foregoing is presented more in the spirit of illustration of our ideas than as a rigorous treatment of a theoretical model. We think of our method as a first stab at the problem of providing a robust but cointegration-free method for estimating dynamic wealth effects, and we hope that more rigorous modeling and structural estimation will follow. But we anticipate that such approaches will confirm the basic dynamics captured by our method, in part because we think those dynamics have been reflected in the results obtained in most

¹³See Carroll and Slacalek (2006) for arguments that microeconomic data suggest sticky expectations rather than habits are the right explanation.

¹⁴We take the expectation of lagged consumption growth because in a small sample with only two shocks, the coefficient estimates from an ordinary least squares (OLS) regression would be severely biased by the tendency of OLS to want to fit the ex-post experience. If we had a large number of shocks in our example, this problem would be smaller, but we wanted to keep the example as simple and ‘toylike’ as possible to aid comprehension. And since our empirical work below will regress consumption on the expectation of lagged consumption, using $\mathbb{E}_{t-2} \Delta \mathbf{C}_{t-1}$ increases the coherence between our exercises on simulated and empirical data.

of the recent literature on structural estimation of more complicated macroeconomic models, which invariably find a strong component of ‘habit formation.’

Our method also has the advantage that it allows for a transparent generalization for comparing the effects of shocks to different kinds of wealth. If housing wealth is measured by \mathbf{B}_t^h and financial wealth by \mathbf{B}_t^f , our method boils down to estimating

$$\mathbf{C}_{t+n} - \mathbf{C}_t = (\mathbf{B}_{t+n}^f - \mathbf{B}_t^f)\kappa_f + (\mathbf{B}_{t+n}^h - \mathbf{B}_t^h)\kappa_h$$

which is sufficiently simple that the respective κ ’s might almost serve as definitions rather than estimates of the sizes of the respective wealth effects at the defined horizon n .¹⁵

3 Estimates Based on Consumption Growth Dynamics

Our estimation approach exploits the robust empirical fact that aggregate consumption growth responds only sluggishly to shocks. The most persuasive evidence that such sluggishness exists is the reluctant introduction of habits into quantitative macroeconomic models in the last few years, despite the evident distaste for the habit formation assumption on the part of many researchers. Models that include habits are proliferating because they can match the core empirical fact of sluggish consumption growth along with attendant implications for asset pricing and other empirical phenomena.

In implementing our method in actual data, the first step is to estimate the degree of stickiness in consumption growth in (3). But there is a problem: The producer of the consumption data documents a variety of sources of measurement error in that data (Bureau of Economic Analysis, 2006). Furthermore, anyone who has been involved in real-time consumption forecasting knows that there are large transitory elements of spending (e.g. hurricane-related purchases) that are not incorporated in the theory that leads to (3).

Fortunately, these problems can be largely overcome when χ is estimated with instrumental variables estimation using instruments dated $t - 2$ or earlier.¹⁶ These estimates suggest a serial correlation coefficient for ‘true’ consumption growth in the neighborhood of 0.7 (whether the measure of spending is total consumption expenditures, spending on nondurables and services, or spending on nondurables

¹⁵As noted before, a reasonable objection to this is that movements in \mathbf{B}^f and \mathbf{B}^h are not likely to be exogenous in the econometrically required sense. Hence the need for more sophisticated models.

¹⁶As shown by Muellbauer (1988) and Sommer (2007), in a simple habit formation model (like the one sketched above), time aggregation causes a moving average MA(2) process in consumption growth. But because the MA(2) coefficient is generally small, using instruments as of time $t - 2$ induces essentially no bias, as illustrated in Table 5 below.

alone). The evidence below confirms that this finding holds robustly for alternative sets and alternative lags of instrumental variables.

3.1 Estimating the Wealth Effect

To estimate the wealth effect, we must modify Sommer’s methodology in several directions. First, the ultimate goal here is to obtain an estimate of the marginal propensity to consume out of wealth. But (3) is written in terms of the growth rate of consumption. Even if the model were estimated as a just-identified system where the only instrument for lagged consumption growth is lagged changes in wealth, the result would be a relationship between the *growth rate* of wealth and the *growth rate* of consumption, which is not an MPC. Worse, this approach makes no sense if wealth is split up into a housing and a financial component. If the null hypothesis is that the MPCs out of the two components are equal then the coefficients on their log changes will *not* be identical unless financial and housing wealth are the same size in every period (in which case their differential effects would not be identified!).

There is a simple solution to these problems, which is to use the ratio of changes in wealth to an initial level of consumption rather than wealth growth.¹⁷ That is, if we define

$$\begin{aligned}\partial\mathbf{C}_t &= (\mathbf{C}_t - \mathbf{C}_{t-1})/\mathbf{C}_{t-5} \\ \partial\mathbf{B}_{t-1} &= (\mathbf{B}_{t-1} - \mathbf{B}_{t-2})/\mathbf{C}_{t-5}\end{aligned}$$

and so on, then a first-stage regression of the form

$$\partial\mathbf{C}_t = \alpha_0 + \alpha_1\partial\mathbf{B}_{t-1} \tag{4}$$

yields a direct estimate of the marginal propensity to consume in quarter t out of a change in wealth in quarter $t - 1$. Furthermore, if \mathbf{B}^f and \mathbf{B}^h are the financial and housing components of wealth, a first-stage regression of the form

$$\partial\mathbf{C}_t = \alpha_0 + \alpha_1\partial\mathbf{B}_{t-1}^f + \alpha_2\partial\mathbf{B}_{t-1}^h \tag{5}$$

yields directly comparable estimates of relative MPCs.

The reader may wonder why the wealth variables in (5) are lagged one period. This is for several reasons. First, wealth in our source (the Flow of Funds Accounts) is measured at a point in time (on the last day of the quarter), while consumption occurs continuously throughout a quarter. If we were to use a measure of wealth with the same time subscript as our measure of consumption, in practice that would be incorporating information that was revealed to the consumer only late in the quarter as though the consumer could have known about it early in the quarter. Second, there is a potentially serious simultaneity problem with looking at the relationship

¹⁷Because we will later be using variables with lags up to a year, the ‘initial’ level here is defined as consumption five quarters before the current quarter.

between current consumption and current wealth: Maybe innovations to both are driven by some exogenous unmeasured third variable (growth expectations, say). Then if asset markets respond instantly to new information (as they should to prevent arbitrage; the random walk proposition is much closer to holding true for asset prices than for consumption), the coefficient on wealth would reflect some of this simultaneity bias rather than a ‘pure’ marginal propensity to consume. Finally, the most useful context in which empirical work like this might be performed is in forecasting high frequency consumption movements. To do that, one needs to have lagged, not contemporaneous, variables on the right hand side.

Regressions of the form (4) or (5) pass all the standard tests of instrument validity and therefore justify estimation of an IV equation of the form

$$\partial \mathbf{C}_t = \gamma + \chi \partial \mathbf{C}_{t-1} + \varepsilon_t \quad (6)$$

where γ is an unimportant constant.

Given an initial (current-quarter) MPC out of wealth of $\underline{\kappa}$ and a serial correlation coefficient χ for $\partial \mathbf{C}$, the usual infinite horizon formula implies that the ultimate effect on the level of consumption (the ‘eventual MPC’) from a unit innovation to wealth is¹⁸

$$\bar{\kappa} = \frac{\underline{\kappa}}{1 - \chi}.$$

Our interpretation of the econometric object we call the ‘eventual MPC’ is that it really reflects the medium-run dynamics of consumption (over the course of a few years); that is, the effects over a time frame short enough that the consequences of the consumption decisions have not had time to have a substantial impact on the level of wealth and to induce general equilibrium offsets. Thus the distinction between what we are calling the ‘eventual’ MPC and what comes out of a cointegration analysis is that in principle the cointegration analysis characterizes some average characteristics of the whole 45-year sample, while our results reflect average dynamics over a much shorter horizon.

Returning to the main thrust, the simplest way to estimate the “eventual MPC” would have been to directly report the relevant coefficient estimates on one-quarter-lagged $\partial \mathbf{B}$ from the first-stage regressions. If that MPC had been α then the fact that $\alpha = \chi \underline{\kappa}$ implies that the eventual MPC could have been estimated from

$$\bar{\kappa} = \frac{\alpha}{\chi(1 - \chi)},$$

¹⁸As a digression, this seems a good place to explain why we introduce the term ‘eventual’ MPC; an earlier draft called this object the ‘long-run MPC.’ We are sympathetic to the objection that in a general equilibrium context it is not clear what ‘long-run MPC out of wealth’ means, because in the long run the amount of wealth is endogenous with respect to consumption choices; indeed, one interpretation of the cointegration discussion above is that the only sensible definition of the ‘long-run MPC out of wealth’ is that it is zero. Effectively, we are assuming that our estimates are identified by high- and medium-frequency variation that is largely uncontaminated by the very long run general equilibrium effects that plague cointegration analysis.

where the χ in the denominator adjusts for the fact that the estimated coefficient is on once-lagged rather than the current change in wealth.

However, the coefficient estimates when only a single lag of each of the two measures of wealth was included in the regression were a bit too sensitive to the inclusion of other instruments for us to be comfortable relying upon them directly.¹⁹ However, if the model of serial correlation in true consumption growth is right, it is easy to make an alternative measure of the change in wealth that should capture the relevant facts. For a given value of χ , assuming independent shocks to wealth from quarter to quarter we should have:

$$\Delta \mathbf{C}_t \approx \underline{\kappa} \chi (\Delta \mathbf{B}_{t-1} + \chi \Delta \mathbf{B}_{t-2} + \chi^2 \Delta \mathbf{B}_{t-3} + \chi^3 \Delta \mathbf{B}_{t-4}) + \varepsilon_t.$$

Now define

$$\bar{\partial} \mathbf{B}_t = (\Delta \mathbf{B}_t + \chi \Delta \mathbf{B}_{t-1} + \chi^2 \Delta \mathbf{B}_{t-2} + \chi^3 \Delta \mathbf{B}_{t-3}) / \mathbf{C}_{t-4} \quad (7)$$

and since similarly $\partial \mathbf{C}_t = (\mathbf{C}_t - \mathbf{C}_{t-1}) / \mathbf{C}_{t-5}$ this leads to an approximate equation for $\partial \mathbf{C}$ and $\bar{\partial} \mathbf{B}$ of the form

$$\partial \mathbf{C}_t = \gamma + \alpha \bar{\partial} \mathbf{B}_{t-1}. \quad (8)$$

Under the assumption that the dynamic model of consumption is right, the coefficient estimate on $\bar{\partial} \mathbf{B}_t$ should be the immediate (first-quarter) MPC out of an innovation to wealth.

Thus, the estimate of the eventual MPC out of wealth reported in table 4 is given by

$$\bar{\kappa}_j = \frac{\alpha_j}{\chi(1-\chi)}. \quad (9)$$

for the α_j , $j \in \{f, h\}$ corresponding to the respective measure of wealth.

To summarize, for each of the instrument sets, the procedure is as follows:

1. Estimate (6) by IV, generating the estimate of χ reported in table 4.
2. Construct the estimate of $\bar{\partial} \mathbf{B}$ as per (7).
3. Estimate (8) or the corresponding equation for the other instrument sets, yielding the estimate of the immediate MPC contained in table 3.
4. Construct the estimate of the eventual MPC for table 4 via (9).

The logic of the foregoing is admittedly a bit circular, but the circularity is motivated more by presentational issues than substance: It seemed essential, for streamlined exposition, to be able to report a single statistic as the immediate MPC and a single statistic as the eventual MPC out of wealth shocks. However, when only

¹⁹The estimates are not enormously sensitive—they typically imply eventual MPCs between 0.02 and 0.1.

a single lag of wealth is used in the first-stage regression the coefficient estimates are implausibly sensitive to the exact specification and exactly which instruments are included. When a few lags are used, the sum of the coefficients on the lags tends to yield similar immediate coefficients, but is harder to summarize. Hence the compromise represented by table 3.

3.2 Estimation Results

As a baseline, the first row of table 3 presents the estimation results of the regression (8) of the change in consumption $\partial\mathbf{C}_t$ on a weighted average of the change in wealth over the prior year $\bar{\partial}\mathbf{B}_{t-1}$. Thus, the regression coefficients are now interpretable as the marginal propensity to consume out of changes in wealth in the previous quarter. The reported results are for total personal consumption expenditures (PCE), because the focus here is on the effects of wealth on aggregate demand, but appropriately scaled-down results can be obtained for spending excluding durables, or excluding both durables and services.

The coefficient estimate in this baseline model implies that if wealth grew by \$1 last quarter, then consumption will grow by about \$0.017 more in the current quarter than if wealth had been flat. While this wealth effect is highly statistically robust, lagged wealth growth alone explains only about 14 percent of quarterly consumption growth (as implied by the \bar{R}^2 from the regression of $\partial\mathbf{C}_t$ on a constant and $\bar{\partial}\mathbf{B}_{t-1}, \dots, \bar{\partial}\mathbf{B}_{t-4}$ not reported in table 3).²⁰

The next step is to find a parsimonious set of additional variables that have significant predictive power for consumption growth. There is a traditional set of variables often used in this literature, dating back to the work of Campbell and Mankiw (1989), including the recent performance of stock prices as well as lagged interest rates and income growth rates. However, for our purposes an adequate representation is obtained by augmenting lagged wealth with just two explanatory variables: Lagged unemployment expectations from the University of Michigan’s consumer sentiment survey (to capture changes in economic uncertainty), and the lagged Fed funds rate, which is included in the hope that it will capture some of the effects of monetary policy, leaving the housing wealth variable to capture more exogenous movements in house prices.

The second row shows that when the extra variables are added, the coefficient on the change in wealth is diminished (by about half). This makes sense because the extra variables are correlated with the change in wealth. However, the extra variables also have considerable independent predictive power for consumption growth. Overall, the explanatory power of the regression including both extra measures is almost double the power of the regression that only includes lagged wealth.

²⁰This does not merely reflect time aggregation; even twice-lagged wealth changes have highly statistically significant predictive power for consumption growth.

The third row regresses the consumption change on the change in housing and financial wealth separately; the point estimate of the effect of housing wealth is more than twice as large as the coefficient on financial wealth (which is close to the original estimate of the effect of total wealth). However, the coefficient on housing wealth is much less precisely estimated than the coefficient on financial wealth, and a statistical test indicates that the hypothesis that the two coefficients are actually equal cannot be rejected at the 95 percent significance level. One reason the coefficient on housing wealth is harder to pin down is that housing wealth varies considerably less than financial wealth, as shown in figure 2.

The final row presents our preferred specification, in which financial and housing wealth effects are examined separately from the other explanatory variables. Results are broadly what would be expected from the foregoing: Both coefficients are substantially smaller, and the coefficient on housing wealth is about twice as large as that on financial wealth, but the difference between the two coefficients is not statistically significant. The coefficient on housing wealth is different from zero, at the 0.14 percent level.

The results in this table are not the bottom line, because they reflect only the next-quarter effect on consumption growth. To obtain the eventual MPCs, we need to estimate equation (6) and apply formula (9). Results of these calculations are reported in table 4.

The first column shows that all models find a very substantial, and highly statistically significant, amount of momentum (by which we mean an estimate of $\chi > 0$) in consumption growth. Note also that the regressions that include the extra explanatory variables (which had much greater power for consumption growth) find notably higher estimates of momentum. Furthermore, in experiments not reported here (but available in the replication archive), a much more extensive set of instruments was examined. The bottom line is that any instrument set that has a reasonable degree of predictive power for $\partial \mathbf{C}_t$ (e.g., an \bar{R}^2 of 0.1 or more) generates a highly statistically significant estimate of the χ coefficient. Furthermore, the estimate of χ tends to be larger the better is the performance of the first-stage regression.

The last two columns report the estimated eventual MPCs out of financial and housing wealth. When the MPCs are permitted to differ for financial and housing wealth, the higher immediate MPCs out of housing wealth from table 3 translate into higher eventual MPCs here, with the preferred model estimate (the last row) of an eventual MPC out of housing wealth of 9 cents on the dollar.

One intuition for why the MPC out of financial wealth is substantially lower than that out of housing wealth is evident in figure 2. Financial wealth is considerably more volatile than housing wealth. If the model is really true, these high frequency fluctuations should have considerable power in explaining subsequent spending patterns. In practice, high frequency stock market fluctuations do not seem

to translate into very large subsequent consumption fluctuations, so the coefficient is not estimated to be very large.²¹

3.3 Comparison with Existing Empirical Work

The work most closely related to ours is Case, Quigley, and Shiller (2003) (henceforth CQS), which provides estimates from both a panel of developed countries (since 1975) and a panel of states within the U.S. Using annual data, CQS find a highly statistically significant estimate of the MPC out of housing wealth in the U.S. of around 0.03–0.04. In contrast, the CQS estimate of the MPC out of stock market wealth is small and statistically insignificant. The coefficient on housing wealth is estimated to be highly statistically significantly larger than the coefficient on financial wealth.

But the literature does not speak with one voice. A study by Ludwig and Slok (2004) estimates a larger effect of financial wealth than housing wealth in a panel of 16 OECD countries, and also reports some evidence of an increase in wealth effects over time. Girouard and Blöndal (2001) fail to find consistent results across countries: In some, the housing wealth effect is stronger, while in others the financial wealth effect is stronger (and in some neither was significant). And a study by Dvornak and Kohler (2003) modelled closely on the CQS study but using Australian state-level data finds a larger financial wealth effect than housing wealth effect.

It should be admitted that there are good reasons to be skeptical of results based on macroeconomic or regional data (including our own). Foremost among these is the previously-acknowledged point that movements in asset prices are not exogenous fluctuations; they should be affected by many of the same factors that affect consumption decisions, most notably overall macroeconomic prospects. House prices should depend, in part, on the overall future purchasing power of current and future homeowners, while stock prices should reflect expectations for corporate profits, which are of course closely tied to the broader economy. John Muellbauer and various co-authors (Aron and Muellbauer, 2006, Muellbauer, 2007, Aron, Duca, Muellbauer, Murata, and Murphy, 2008) (using Japanese, South African, U.K. and U.S. data) have attempted to address this problem by including control variables for credit market liberalizations and other time varying conditions. But to isolate a ‘pure’ housing wealth effect, one would want data on spending by individual households before and after some truly exogenous change in their house values, caused for example by the unexpected discovery of neighborhood sources of pollution.

The perfect experiment observed in the perfect microeconomic dataset is not available. Many authors have attempted to measure housing wealth effects using

²¹Figure 2 actually shows levels of wealth rather than differences $\bar{\delta}\mathbf{B}$ used in the regressions. However, the volatility in levels is transferred into differences. Consequently, it turns out that the standard deviation of the financial wealth measure $\bar{\delta}\mathbf{B}^f$ is about three times as large as that of $\bar{\delta}\mathbf{B}^h$.

microeconomic datasets, but heroic assumptions usually must be made in order to produce estimates, because the existing datasets were not designed with this question in mind.

Given these problems, it is not surprising that the results from microeconomic studies are even more heterogeneous than those from macroeconomic data.

Recent studies by Attanasio, Blow, Hamilton, and Leicester (2008), and Campbell and Cocco (2006) represent both the wide spectrum of views and the best available microeconomic evidence and methodologies.

Disney, Gathergood, and Henley (2008) find an MPC out of unanticipated shocks to housing wealth of only 0.01, after controlling for expectations of future financial conditions. They show that without such controls, the estimated MPC is considerably higher, a result that strongly suggests that the macroeconomic correlation evident in both U.K. and U.S. data reflects causality from general economic conditions to both consumption and asset prices, rather than a direct housing wealth effect.

On the other hand, Campbell and Cocco (2006) also use British data (this time, from the U.K. Family Expenditure Survey and from regional house price surveys), but find a large housing wealth effect, which is different for young and old households; they find a statistically significant *elasticity* of consumption to house prices of about 1.7 among older homeowners, but no significant effect among young renters.

Attanasio, Blow, Hamilton, and Leicester (2008), in contrast, find that consumption of young renters *is* positively associated with house price changes, which again suggests that both consumption and house prices are responding to an unobserved aggregate. Additional microeconomic estimates of the wealth effect are reported in Engelhardt (1996), Juster, Lupton, Smith, and Stafford (2001), Lehnert (2003), Levin (1998) and Bostic, Gabriel, and Painter (2005).

Stepping back from the conflicting details of the disparate studies, perhaps the most useful observation is that even if it is true that the ‘pure’ housing wealth effect is modest, if a macroeconomic policymaker wants to know what to expect for future consumption growth given a particular recent path of aggregate wealth shocks, it may matter more whether the forecast is reliable than whether the mechanism is a direct wealth effect, a reflection of an omitted variable like growth expectations, or a reflection of a difficult-to-measure variable like credit conditions. If, for example, a collapse in house prices properly signals a collapse in consumption, the precise mechanism by which consumption will collapse may not be so important.

3.4 The Relevance of Various Wealth Effect Channels

Despite the obvious limitations of aggregate data, we now attempt to decompose the total response of spending to wealth into the parts due to the five channels outlined

above in section 2:²² 1. The ‘statistical’ effect because the stream of housing services is included in total PCE and depends on housing wealth, 2. The possibility that the MPC out of a particular kind of wealth might depend on its degree of liquidity, 3. Collateral constraints might be important; and 4. The cross-sectional distribution of wealth might matter.

To address the relevance of the statistical effect, we have re-estimated the model measuring consumption with total PCE *excluding* housing services. The results are in line with our baseline: The housing wealth effect ($\bar{\kappa}_h = 0.070$) remains highly statistically significant and roughly twice as large as the financial wealth effect ($\bar{\kappa}_f = 0.039$). As a caveat it is worth mentioning that this alternative specification addresses the problem only when utility is additively separable in housing services and the rest of PCE.²³

The second and third channels are difficult to assess separately and are both driven by financial innovation. Iacoviello and Neri (2007) argue that the recent increase in liquidity of housing is captured in the higher loan-to-value (LTV) ratio.²⁴ In addition, as pointed out by Muellbauer (2007), the rise in LTV ratios (and the reduction in down-payments) increases consumption of young credit-constrained first-time home buyers. On the other hand, the falling relevance of credit constraints (both in terms of the number of households they affect and their extent) has likely weakened the wealth effect.²⁵

Muellbauer (2007) constructs an indicator of credit market conditions based on the Federal Reserve’s Senior Loan Officer Survey question about the willingness of banks to make consumer installment loans (see the installment loans credit indicator in Figure 4 of his paper). Possibly because of the deregulation and restructuring of the U.S. housing finance system (see e.g., McCarthy and Peach, 2002), the indicator rose markedly around 1984, a movement which likely drives much of the significant increase in the housing wealth effect reported by Muellbauer (2007). The split-sample regressions (pre-1985 and post-1984) we have estimated with our method confirm this finding: The eventual housing wealth effect rose from only 0.03 to 0.12 (while the financial wealth effect actually fell from 0.08 to 0.03).²⁶ Much of the recent

²²Detailed results reported in this section are available in the replication archive.

²³The literature (e.g., Davis and Martin, 2005 and Piazzesi, Schneider, and Tuzel, 2007) seems to agree that the complementarities between non-housing consumption and housing services are modest.

²⁴Muellbauer (2007), Figure 4, plots the data from the American Housing Survey, in which the LTV increased from roughly 90 percent in 2000 to about 95 percent in 2005. In a different dataset from the Federal Housing Finance Board (used, e.g., by Iacoviello and Neri, 2007), the LTV ratio rose by a few percentage points after 1995. Dynan and Kohn (2007), p. 18, report that in the U.S. Survey the share of households with some debt increased from 70 percent in 1983 to 77 percent in 2004.

²⁵Comparing pre-1983 and post-1989 U.S. data Iacoviello and Neri (2007) estimate that the income share of credit-constrained consumers fell from 0.32 to 0.20.

²⁶Iacoviello and Neri (2007) also report that their estimated dynamic stochastic general equilibrium model with housing sector suggests that the response of consumption to shocks rose considerably (after 1988). Slacalek (2009) reports that the wealth effect increased (statistically significantly) after 1989 in many of the 16 industrialized countries he investigates.

literature thus seems to agree that the impact of housing wealth on consumption has been rising in a period (post-1985 or so) which coincides with the intense financial innovation. This evidence is suggestive of a potential causal link. (Of course, as in many other applications, econometric methods like ours do not make it possible to make a final conclusion on the direction of causality.) Our split-sample regressions thus point to a substantial role of financial innovation (channels number 2 and 3) in determining the size of the housing wealth effect. While there are many distinct ways in which financial markets affect the transmission between wealth and consumption, on balance it does seem likely that financial innovation may have made consumption more responsive to housing wealth shocks.

The fourth channel that might affect the size of the wealth effect on aggregate level is the cross-sectional distribution of various classes of assets. Estimates with aggregate data implicitly identify the marginal propensity to consume out of wealth averaged across households: $\bar{\kappa} = (1/N) \sum_{i=1}^N \kappa_i(\mathbf{B}_i)\omega_i$, where the marginal propensity κ_i ²⁷ of each consumer decreases with his wealth \mathbf{B}_i (due to the diminishing role of the precautionary saving motive), and ω_i is the household's weight in the aggregate statistic. Housing is considerably more evenly distributed than financial assets: In the U.S. Survey of Consumer Finances (SCF) of 2004 the top five percent of households (by net worth) held 26.3 percent of the total value of houses (or \$5.0 trillion) but 57.9 percent of financial assets (or \$12.2 trillion) (see Kennickell, 2006, Table 11a). Unfortunately, it is difficult to assess quantitatively by how much the aggregate MPC out of housing wealth would fall if we exogenously imposed that housing has the same distribution as financial assets.²⁸ However, it is well-known that both housing and financial wealth of the richest households has since 1995 grown very rapidly (see, e.g., Survey of Consumer Finances, 2007).²⁹ This shift has probably, if anything, weakened wealth effects. However, the change seems likely to be modest because theory suggests that the spending of the rich people should not react much to shocks (both because of the weak precautionary saving motive and the irrelevance of liquidity constraints).

3.5 Alternative Specifications

Table 5 demonstrates the robustness of our estimates of the wealth effects to three alternative specifications of the model. The top panel considers an alternative instrument set for lagged consumption growth \mathbf{C}_{t-1} in (6), which consists of the growth

²⁷The function $\kappa_i(\cdot)$ differs across households because of, e.g., demographics.

²⁸In theoretical models the marginal propensity is determined by the coefficient of relative risk aversion, the discount rate, and all the other parameters of the model. The key problem in calculating counterfactual implications like this is that the standard consumption models succeed in matching neither the upper tail of the distribution of net worth nor its composition.

²⁹Following the burst of the internet bubble, financial asset of most people fell somewhat between the 2001 and 2004 waves of SCF.

rate of stock prices, change in unemployment rate, the growth rate of disposable income and the interest rate spread. The second panel investigates the robustness of estimates of χ , α and $\bar{\kappa}$ to the inclusion of only lags $t - 3$ and $t - 4$ of these instruments, a procedure which is an appropriate method under MA(2) disturbances but the instruments have lower forecasting power for consumption growth than the baseline method. The third panel shows the estimates from the following *iterative* procedure, which tests how sensitive the estimates of consumption sluggishness χ are to the wealth variable $\bar{\partial}\mathbf{B}_t$. The procedure consists of re-estimating for the second time the IV regression (6) with $\bar{\partial}\mathbf{B}_t$ among instruments instead of $\partial\mathbf{B}_t$ and backing out the estimates of the wealth effect using the updated series for $\bar{\partial}\mathbf{B}_t$, which is calculated using the second-round estimate of χ . Finally, the bottom panel shows the estimates of housing and financial wealth effects implied by a model in which household wealth is split into the two components as follows: net housing wealth is measured as real estate held by households minus mortgages; net financial wealth is measured as total assets net of real estate held by households and non-mortgage liabilities.³⁰

The results suggest that the estimates of consumption sluggishness χ typically lie around 0.6–0.7 and the estimates of the immediate and eventual marginal propensity to consume out of wealth are roughly 0.010 and 0.05, respectively. In addition, the housing wealth effects are consistently larger than the financial wealth effects and are broadly in line with our baseline estimate of 0.09. The methods also achieve better first-stage fit (higher \bar{R}^2) because they are based on a larger set of (valid) instruments than the baseline estimates.

4 Conclusion

Our results suggest that, in U.S. historical experience, housing price movements have typically been associated with substantial subsequent movements in consumer

³⁰The specification is motivated by Calomiris, Longhofer, and Miles (2009), p. 17, who argue that our measure of housing wealth is problematic because “more than half of ‘housing wealth’ consists of non-housing wealth” and claim that this fact biases upward our estimates of the housing wealth effect; they propose an alternative measure which they claim is a better measure of housing wealth, and show that using their measure the housing wealth effect is smaller. But their measure, oddly, subtracts mortgage debt from financial assets to yield a supposed measure of ‘financial wealth,’ while counting housing assets as though they were unencumbered by debt. We considered several measures of housing versus nonhousing wealth in earlier drafts of this paper, but confess that this surprising configuration did not occur to us. Most choices that we did try yielded results similar to or stronger than those we present as our baseline. For example, our estimate of the immediate MPC out of housing wealth of 0.022 for model 7 suggests that the results are quite robust to measuring wealth. (The eventual housing MPC is larger than for the baseline specification because of the high χ for model 7. We see no reason to change our original view that, on balance, the evidence is mildly supportive of a larger effect for housing wealth than for nonhousing wealth, but that the hypothesis that the two wealth effects are equal cannot be rejected. And, as noted at the outset of the paper, we are sympathetic to the possibility that what this literature calls a ‘wealth effect’ may instead be a reflection of a correlation between wealth and other variables that are harder to measure, like growth expectations or credit conditions. Those are questions unlikely to be answerable using aggregate data, though microeconomic and regional data offer hope of eventual resolution of the problem.)

spending. The immediate (first-quarter) impact is estimated to have been relatively small (the immediate quarterly MPC in our preferred model is about 2 cents on the dollar), but over a time span of several years we estimate that it has on average accumulated to the 4–10 cent range. These figures are consistent with evidence from other studies and the experience across U.S. states. Whether the housing wealth effect is substantially larger than the financial wealth effect is more uncertain; while the bulk of the literature seems to point in that direction, in our estimates the size of the differences is not large enough to yield confidence in the conclusion.

For monetary policy purposes, these results suggest that it would be wise for policymakers to keep a close eye on developments in housing markets separately from equity markets, since even the possibility of a significantly higher MPC out of housing wealth can shift the balance of risks in a macroeconomic forecast. Such a perspective, for example, could have helped in understanding and interpreting the surprising strength of the U.S. consumption and residential investment spending in the early 2000s even as the stock market suffered a historic decline.

References

- ARON, JANINE, JOHN DUCA, JOHN MUELLBAUER, KEIKO MURATA, AND ANTHONY MURPHY (2008): “Credit, Housing, Collateral and Consumption: Evidence from the UK, Japan and the US,” mimeo, University of Oxford.
- ARON, JANINE, AND JOHN MUELLBAUER (2006): “Housing Wealth, Credit Conditions and Consumption,” working paper 9, Centre for the Study of African Economies, Oxford University.
- ATTANASIO, ORAZIO, LAURA BLOW, ROBERT HAMILTON, AND ANDREW LEICESTER (2008): “Consumption, House Prices, and Expectations,” *Bank of England Working Paper Number 271*.
- BERTAUT, CAROL (2002): “Equity Prices, Household Wealth, and Consumption Growth in Foreign Industrial Countries: Wealth Effects in the 1990s,” International Finance Discussion Paper 724, Federal Reserve Board.
- BOSTIC, RAPHAEL, STUART GABRIEL, AND GARY PAINTER (2005): “Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data,” mimeo, University of Southern California.
- BUREAU OF ECONOMIC ANALYSIS (2006): “Updated Summary NIPA Methodologies,” Survey of Current Business, November, Bureau of Economic Analysis, available at http://www.bea.gov/scb/pdf/2006/11November/1106_nipa_method.pdf.
- BYRNE, JOSEPH P., AND E. PHILLIP DAVIS (2003): “Disaggregate Wealth and Aggregate Consumption: An Investigation of Empirical Relationships for the G7,” *Oxford Bulletin of Economics and Statistics*, 65(2), 197–220.
- CALOMIRIS, CHARLES, STANLEY D. LONGHOFFER, AND WILLIAM MILES (2009): “The (Mythical?) Housing Wealth Effect,” Working Paper 15075, NBER.
- CAMPBELL, JOHN Y., AND JOÃO F. COCCO (2006): “How Do House Prices Affect Consumption? Evidence from Micro Data,” *Journal of Monetary Economics*, forthcoming.
- CAMPBELL, JOHN Y., AND ANGUS S. DEATON (1989): “Why Is Consumption So Smooth?,” *Review of Economic Studies*, 56, 357–74.
- CAMPBELL, JOHN Y., AND N. GREGORY MANKIW (1989): “Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence,” in *NBER Macroeconomics Annual*, ed. by Olivier J. Blanchard, and Stanley Fischer, Cambridge, MA. MIT Press.

- CARROLL, CHRISTOPHER D. (2004): “Housing Wealth and Consumption Expenditure,” *Paper Prepared for Academic Consultants Meeting of Federal Reserve Board, January 2004*, At <http://econ.jhu.edu/people/ccarroll/papers/FedHouseWealthv2.pdf>.
- CARROLL, CHRISTOPHER D., AND JIRI SLACALEK (2006): “Sticky Expectations and Consumption Dynamics,” mimeo, Johns Hopkins University.
- CASE, KARL E., JOHN M. QUIGLEY, AND ROBERT J. SHILLER (2003): “Comparing Wealth Effects: The Stock Market Versus the Housing Market,” *Institute of Business and Economic Research Working Paper No. W01-004*.
- CATTE, PIETRO, NATHALIE GIROUARD, ROBERT PRICE, AND CHRISTOPHE ANDRE (2004): “Housing Markets, Wealth and the Business Cycle,” OECD Economics Department Working Papers 394, OECD.
- CHETTY, RAJ, AND ADAM SZEIDL (2005): “Consumption Commitments: Neoclassical Foundations for Habit Formation,” mimeo, University of California, Berkeley.
- 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*.
- DAVIS, MORRIS A., AND ROBERT F. MARTIN (2005): “Housing, House Prices, and the Equity Premium Puzzle,” Finance and Economics Discussion Series 13, Board of Governors of the Federal Reserve System.
- DISNEY, RICHARD, JOHN GATHERGOOD, AND ANDREW HENLEY (2008): “House Price Shocks, Negative Equity, and Household Consumption in the United Kingdom,” *Manuscript, University of Nottingham*.
- DVORNAK, NIKOLA, AND MARION KOHLER (2003): “Housing Wealth, Stock Market Wealth and Consumption: A Panel Analysis for Australia,” *Reserve Bank of Australia Research Discussion Paper 23*.
- DYNAN, KAREN E. (2000): “Habit Formation in Consumer Preferences: Evidence from Panel Data,” *American Economic Review*, 90(3).
- DYNAN, KAREN E., AND DONALD L. KOHN (2007): “The Rise in U.S. Household Indebtedness: Causes and Consequences,” International Finance Discussion Paper 37, Board of Governors of the Federal Reserve System.
- ENGELHARDT, GARY V. (1996): “House Prices and Home Owner Saving Behavior,” *Regional Science and Urban Economics*, 26, 313–336.

- FERNANDEZ-CORUGEDO, EMILIO, SIMON PRICE, AND ANDREW BLAKE (2003): “The Dynamics of Consumers’ Expenditure: The UK Consumption ECM Redux,” working paper 204, Bank of England.
- FLAVIN, MARJORIE A. (1981): “The Adjustment of Consumption to Changing Expectations About Future Income,” *Journal of Political Economy*, 89(5), 974–1009.
- FUHRER, JEFFREY C. (2000): “An Optimizing Model for Monetary Policy: Can Habit Formation Help?,” *American Economic Review*, 90(3).
- GIROUARD, NATHALIE, AND SVEINBJÖRN BLÖNDAL (2001): “House Prices and Economic Activity,” *OECD Economic Department Working Paper 279*.
- HAMBURG, BRITTA, MATHIAS HOFFMANN, AND JOACHIM KELLER (2005): “Consumption, Wealth and Business Cycles in Germany,” discussion paper 16, Deutsche Bundesbank.
- IACOVIELLO, MATTEO, AND STEFANO NERI (2007): “Housing Market Spillovers: Evidence from an Estimated DSGE Model,” mimeo, Boston College.
- JUSTER, F. THOMAS, JOSEPH LUPTON, JAMES P. SMITH, AND FRANK STAFFORD (2001): “The Decline in Household Saving and Wealth Effect,” *Mimeo, University of Michigan*.
- KENNICKELL, ARTHUR B. (2006): “Currents and Undercurrents: Changes in the Distribution of Wealth 1989–2004,” Survey of Consumer Finances working paper, Board of Governors of the Federal Reserve System, Available at <http://www.federalreserve.gov/pubs/oss/oss2/method.html>.
- KHALIFA, SHERIF (2004): “Brief Summaries of Principal Papers on Housing Wealth Effects,” *Manuscript, Johns Hopkins University*, Available at <http://econ.jhu.edu/people/ccarroll/wealthLit.html>.
- LEHNERT, ANDREAS (2003): “Housing, Consumption, and Credit Constraints,” *Manuscript, Board of Governors of the Federal Reserve System*.
- 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.
- LEVIN, LAURENCE (1998): “Are Assets Fungible? Testing the Behavioral Theory of Life Cycle Savings,” *Journal of Economic Behavior and Organization*, 36, 59–83.

- LUDWIG, ALEXANDER, AND TORSTEN SLOK (2004): “The Relationship between Stock Prices, House Prices and Consumption in OECD Countries,” *Topics in Macroeconomics*, 4(1, Article 4), available at <http://ideas.repec.org/a/bep/mactop/v4y2004i1p1>.
- 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.
- MUELLBAUER, JOHN (1988): “Habits, Rationality and Myopia in the Life Cycle Consumption Function,” *Annales d’Economie et de Statistique*, 9, 47–70.
- MUELLBAUER, JOHN N. (2007): “Housing, Credit, and Consumer Expenditure,” *Housing, Housing Finance, and Monetary Policy*.
- PIAZESSI, MONIKA, MARTIN SCHNEIDER, AND SELALE TUZEL (2007): “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 83, 531–569.
- PICHETTE, LISE, AND DOMINIQUE TREMBLAY (2003): “Are Wealth Effects Important for Canada?,” working paper 30, Bank of Canada.
- REIS, RICARDO (2006): “Inattentive Consumers,” *Journal of Monetary Economics*, 53(8), 1761–1800.
- RUDD, JEREMY, AND KARL WHELAN (2006): “Empirical Proxies for the Consumption-Wealth Ratio,” *Review of Economic Dynamics*, 9, 34–51.
- SAMUELSON, PAUL A. (1969): “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” *Review of Economics and Statistics*, 51, 239–46.
- SIMS, CHRISTOPHER A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690, available at <http://ideas.repec.org/a/eee/moneco/v50y2003i3p665-690.html>.
- SLACALEK, JIRI (2009): “What Drives Personal Consumption? The Role of Housing and Financial Wealth,” *The B.E. Journal of Macroeconomics*, 9(1), article 37.
- SOMMER, MARTIN (2007): “Habit Formation and Aggregate Consumption Dynamics,” *The B.E. Journal of Macroeconomics—Advances*, 7(1), Article 21.
- SUMMERS, LAWRENCE H. (1981): “Capital Taxation and Accumulation in a Life Cycle Growth Model,” *American Economic Review*, 71(4), 533–544, Available at <http://ideas.repec.org/a/aea/aecrev/v71y1981i4p533-44.html>.

SURVEY OF CONSUMER FINANCES (2007): “2004 SCF Chartbook,” electronic document, Board of Governors of the Federal Reserve System, Available at <http://www.federalreserve.gov/pubs/oss/oss2/2004/scf2004home.html>.

Appendix: Description of Data

Consumption Total personal consumption expenditures; source: National Income and Product Accounts, Bureau of Economic Analysis.

(Total) Wealth Net worth; source: Flow of Funds Accounts, Board of Governors of the Federal Reserve System.

Financial wealth Sum of equity by households, corporate equity by private pension funds, government retirement fund, bank trusts and estates, closed end funds, mutual funds and life insurance companies; source: Flow of Funds Accounts, Board of Governors of the Federal Reserve System.

Housing wealth Net worth minus Financial wealth.


Population Source: National Income and Product Accounts, Bureau of Economic Analysis.

Fed funds rate Source: Fred II database of St. Louis Fed, <http://research.stlouisfed.org/fred2/>.


Unemployment expectations Question 12 of the University of Michigan Survey of Consumer Expectations; source: Survey Research Center, <http://www.sca.isr.umich.edu/>.

Consumption and wealth are measured in real per capita terms, deflated with the consumption deflator. All results are reported for quarterly data, 1960Q1–2007Q4.

Figure 1 Reaction to a Shock to Market Wealth Followed by a Shock to Income Growth under Frictionless and Sticky Expectations



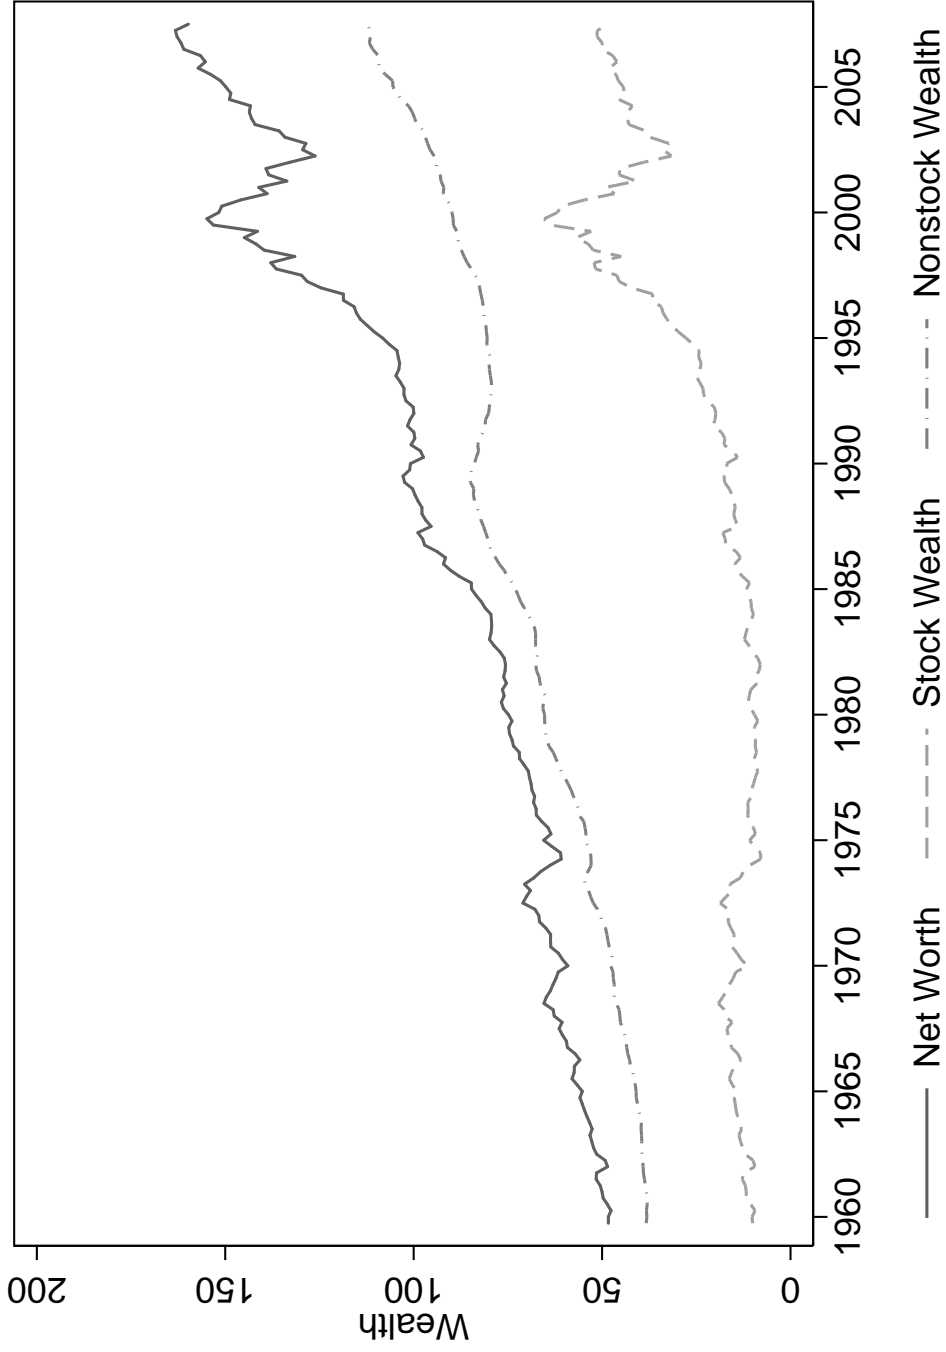
(a) Dynamics of the Consumption Ratio C



(b) Dynamics of the Nonhuman Wealth Ratio B

Note: Both variables normalized with permanent income \mathbf{P}_t . Calibration: $\rho = 2$, $\beta = 1 - \vartheta = 0.99$, $1 + g = G = 1.015^{1/4}$, $1 + r = R = G^\rho/\beta$, $\Pi = 0.25$, wealth shock = 1, income growth shock: $1 + \bar{g} = \bar{G} = 1.025^{1/4}$.

Figure 2 Components of Household Wealth



Note: Per capita real wealth figures in thousands of year 2000 dollars. Net worth is our measure of total wealth.

Table 1 Estimates of the Wealth Effect in Simulated Data—The Cointegration Method, Frictionless Model

$$C_t = \kappa B_t + \varepsilon_t$$

True κ	Estimated κ		
	$n = 20$	$n = 40$	$n = 60$
0.0137	-0.0486	-0.0185	-0.0082

Notes: Both variables normalized with permanent income P_t . “True κ ” = $(R - (R\beta)^{1/\rho})/R$. Calibration: $\rho = 2$, $\beta = 0.99$, $R = G^\rho/\beta$, $\Pi = 0.25$, $\underline{G} = 1.015^{1/4}$, $G = 1.025^{1/4}$, wealth shock = 1.

Table 2 Estimates of the Wealth Effect in Simulated Data—The COS Method

$$\Delta C_t = \chi \mathbb{E}_{t-2} \Delta C_{t-1} + \alpha \Delta B_{t-1} + \varepsilon_t$$

	True	Estimated					
		$n = 20$		$n = 40$		$n = 60$	
χ	$(R - (R\beta)^{1/\rho})/R$	χ	$\alpha/\chi(1 - \chi)$	χ	$\alpha/\chi(1 - \chi)$	χ	$\alpha/\chi(1 - \chi)$
0.75		0.694		0.724		0.734	
0.75	0.0137	0.698	0.0117	0.741	0.0136	0.751	0.0067

Notes: Calibration: $\rho = 2$, $\beta = 0.99$, $R = G^\rho/\beta$, $\Pi = 0.25$, $\underline{G} = 1.015^{1/4}$, $G = 1.025^{1/4}$, wealth shock = 1. ΔC_{t-1} instrumented with ΔC_{t-2} .

Table 3 Immediate Effect of Wealth on Consumption

$$\partial C_t = \alpha_0 + \alpha_1 \bar{\partial B}_{t-1} + \alpha_2 \bar{\partial B}_{t-1}^f + \alpha_3 \bar{\partial B}_{t-1}^h + \alpha_4 MU_{t-1} + \alpha_5 FF_{t-1}$$

Next-Quarter Effect of \$1 Change in Wealth			Extra Variables		Test of $\bar{\partial B}^f = \bar{\partial B}^h$	\bar{R}^2
Total $\bar{\partial B}_{t-1}$	Financial $\bar{\partial B}_{t-1}^f$	Housing $\bar{\partial B}_{t-1}^h$	Unemp Exp MU_{t-1}	Fed Fund FF_{t-1}		
0.017*** (0.004)						0.130
0.009*** (0.003)			0.086*** (0.032)	-0.399* (0.209)		0.222
	0.016*** (0.004)	0.039*** (0.011)			0.066	0.138
	0.008*** (0.003)	0.018** (0.008)	0.082** (0.034)	-0.411* (0.211)	0.271	0.225

Notes: Sample period is 1960Q1–2007Q4. Standard errors in parentheses. {*,**,***}=Statistical significance at {10, 5, 1} percent. Coefficients on wealth variables reflect MPCs in the quarter following a wealth change: For example, the coefficient 0.017 in the first row implies that a one dollar increase in wealth in the previous quarter translates into a 1.7 cent increase in consumption in the current quarter. The wealth variables are from the Flow of Funds balance sheets for the household sector. MU is the fraction of consumers who expect the unemployment rate to decline over the next year minus the fraction who expect it to increase. FF is the nominal Fed funds rate. The wealth and consumption variables were normalized by the level of consumption expenditures at $t-4$ to correct for the long-term trends in consumption and wealth. The equations without the extra variables exhibited serial correlation and so standard errors for those equations are corrected for serial correlation using the Newey–West procedure with 4 lags.

Table 4 Consumption Growth Momentum and the Eventual MPC

$$\partial\mathbf{C}_{t+1} = \mathbf{c}_0 + \chi\mathbb{E}_{t-1}\partial\mathbf{C}_t + \varepsilon_{t+1}$$

Variables used to forecast	Consumption Growth Momentum Coefficient	Implied Eventual MPC out of		
		Total \mathbf{B}	Financial \mathbf{B}^f	Housing \mathbf{B}^h
$\mathbb{E}_{t-1}\partial\mathbf{C}_t$	χ			
\mathbf{B}	0.58** (0.23)	0.070		
$\mathbf{B},$ MU, FF	0.76*** (0.14)	0.048		
$\mathbf{B}^f, \mathbf{B}^h$	0.45** (0.20)		0.064	0.159
$\mathbf{B}^f, \mathbf{B}^h,$ MU, FF	0.71*** (0.13)		0.041	0.087

Notes: Sample period is 1960Q1–2007Q4. Standard errors are in parentheses. $\{*, **, ***\}$ = Statistical significance at $\{10, 5, 1\}$ percent. The eventual MPCs are calculated from the formula $\alpha_j/\chi(1-\chi)$ where α_j is the corresponding next-quarter MPC estimated in table 3. Standard errors for all equations are heteroskedasticity and serial-correlation robust. When more instruments are used to forecast $\partial\mathbf{C}_t$ (for example, interest rate spread and the change in unemployment over the previous year), the estimate of χ tends to rise further and the standard error falls further. The measure of the change in wealth used for the regressions is the $\partial\mathbf{B}$ measure defined in the text, as this can be measured without an estimate of χ , unlike the $\bar{\partial}\mathbf{B}$ measures used in the previous table.

Table 5 Wealth Effect on Consumption—Alternative Specifications

Model	χ	Immediate Effect of \$1 Change in Wealth		Eventual Effect of \$1 Change in Wealth			Test of $\partial B^f = \partial B^h$	\bar{R}^2
		Total ∂B_{t-1}	Financial ∂B_{t-1}^f	Housing ∂B_{t-1}^h	Total ∂B_{t-1}	Financial ∂B_{t-1}^f		
Alternative Instruments								
M1	0.731*** (0.149)	0.013*** (0.005)			0.066			0.301
M2	0.737*** (0.153)	0.011*** (0.004)	0.024** (0.010)		0.055	0.123	0.162	0.309
Instruments as of Time $t - 3$ and $t - 4$								
M3	0.634*** (0.199)	0.015*** (0.006)			0.063			0.301
M4	0.601*** (0.197)	0.013** (0.005)	0.028** (0.012)		0.053	0.116	0.187	0.307
Iterative Method								
M5	0.741*** (0.132)	0.009*** (0.003)			0.047			0.223
M6	0.719*** (0.131)	0.008*** (0.003)	0.018** (0.008)		0.041	0.087	0.273	0.224
Real Estate–Non-Real Estate Split								
M7	0.855*** (0.148)	0.007*** (0.002)	0.022 (0.014)		0.060	0.176	0.309	0.119

Notes: Sample period is 1960Q1–2007Q4. Standard errors in parentheses. {*, **, ***}=Statistical significance at {10,5,1} percent. Models M1 and M2 are estimated with an alternative instrument set which consists of growth rate of stock prices, change in unemployment rate, growth rate of disposable income and interest rate spread. Models M3 and M4 are estimated with lags $t - 3$ and $t - 4$ of these instruments. Models 5 and 6 are estimated using the iterative procedure described in section 3.5.