How Large Is the Housing Wealth Effect? A New Approach *, †

Christopher D. Carroll
ccarroll@jhu.edu
JHU

Misuzu Otsuka
misuzuotsuka@adb.org
ADB

Jirka Slacalek
jslacalek@diw.de
DIW

October 18, 2006

Abstract

This paper presents a simple new method for estimating the size of 'wealth effects' on aggregate consumption. The method exploits the well-documented sluggishness of consumption growth (often interpreted as 'habits' in the asset pricing literature) to distinguish between short-run and long-run 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 long-run effect around 9 cents. Consistent with most recent studies, we find a housing wealth effect that is substantially larger than the stock wealth effect. We believe that our approach has sounder theoretical foundations than the currently popular cointegration-based estimation methods, because neither theory nor evidence provides any reason for faith in the existence of a stable cointegrating vector.

Keywords: housing wealth, wealth effect, consumption dynamics, asset price bubbles

JEL classification: E21, E32, C22

*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 are grateful to Emmanuel de Veirman and Sherif Khalifa for excellent research assistance.

†Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD; Otsuka: Asian Development Bank, Manila, Philippines; Slacalek, German Institute for Economic Research, DIW Berlin, Germany.

We thank Robert Shiller, Robert Hall, Nicholas Souleles, Stephen Zeldes, the staff of the Fed, and other attendees at the 2004 Academic Consultants' meeting for valuable feedback. The data and econometric programs that generated all of the results in this paper are available at http://www.econ.jhu.edu/people/ccarroll/. The views presented in this paper are the authors', and do not necessarily reflect those of the Asian Development Bank and DIW Berlin.

1 Introduction

The exuberant increase in U.S. property prices in the past few years has provided a useful storyline for analysts needing to explain robust household spending despite anemic wages and a lackluster stock market.The challenge, for economists, is to decide whether such explanations are comforting bedtime stories, or a reliable guide to household behavior for the purpose of economic policymaking.

The traditional starting point for quantitative analysis is an off-the-shelf perfect-foresight consumption model that does not distinguish between forms of wealth; an exogenous increment to financial wealth has the same effect as an equal-sized real estate boom. Conventional wisdom puts these effects somewhere in the neighborhood of 3 to 5 cents on the dollar. Many plausible objections can be made to treating property wealth like corporate shares, but taken as a whole it is unclear whether those objections imply a larger wealth effect for stocks or houses; the answer, in the end, must rely on empirical evidence.

Unfortunately, much of the available evidence comes from cointegration estimation methods with dubious theoretical validity and shaky empirical foundations. The theoretical problem is that cointegration models implicitly require the existence of a stable long-run relationship between consumption, labor income, and wealth. Theory implies no such stability, unless every major facet of the economy is perpetually unchanging. Even for the U.S., the 50 year span of available data has seen major changes in taxes, demographics, productivity growth, financial structure, social insurance, and every other aspect of reality incorporated in the theory (and embodied in the cointegrating vector). It is unsurprising, therefore, that empirical tests strongly suggest instability in the cointegrating vectors.

In response, this paper proposes and implements a new methodology for estimating wealth effects, motivated by the large empirical literature documenting substantial sluggishness in aggregate consumption growth.[1]

The methodology can be thought of as a tool for measuring two key aspects of the wealth effect: Speed and strength. Our estimates imply that the immediate effects of housing wealth fluctuations are much smaller than the medium-run effects. In particular, we find that the immediate (next-quarter) marginal propensity to consume from a \$1 change in housing wealth is about 2 cents, with a final long-run effect of 9 cents. Consistent with most other studies, we find a housing wealth effect that is considerably larger than the stock wealth effect.

The large estimated housing wealth effect suggests that policymakers should keep

[1]The 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 or inattentiveness (see e.g. Fuhrer, 2000; Sommer, 2002; Reis, 2004 and Carroll and Slacalek, 2006).

a close eye on movements in housing prices. But the sluggishness of the estimated adjustment processs also suggests that policymakers will typically have plenty of time to see housing market effects as they come through the pipeline, so even if housing wealth effects are large they may not be hard to handle.

2 Estimates Based on Consumption Growth Dynamics

Our approach exploits the robust empirical fact that aggregate consumption growth responds only sluggishly to shocks. Perhaps the most persuasive evidence that such sluggishness is an unavoidable fact is the reluctant introduction of habit formation parameters into quantitative macroeconomic models in the last few years. Habit formation parameters are proliferating because they alllow models to match sluggish consumption growth, unlike the benchmark random walk model of Hall (1978).[2] A canonical model of Muellbauer (1988) proposes a utility function of the form

$$u(C, H) \quad = \quad \frac{(C - \chi H)^{1-\rho}}{1-\rho},$$

where χ is a parameter capturing the importance of habits H to consumption (if $\chi = 0$, habits are irrelevant), and the habit stock evolves according to

$$H_{t+1} \quad = \quad H_t(1 - \lambda) + (C_t - H_t)\lambda.$$

Dynan (2000) shows that in Muellbauer's model the Euler equation can be approximated by

$$\Delta \log C_{t+1} \quad = \quad c_0 + \chi \Delta \log C_t + \varepsilon_{t+1}, \tag{1}$$

so that the importance of habits can be estimated from the serial correlation parameter in consumption growth.[3] The disturbance term ε_{t+1} summarizes the combined effect of all shocks on consumption growth and includes a part due to the wealth shocks $\Delta \log W_{t+1}$.

Suppose that our actual empirical measure of consumption contains either some transitory measurement error, or transitory elements of spending (e.g. hurricanes) that are not incorporated in the theory that leads to (1). The evolution of measured

[2]It is perhaps useful here to reiterate that the empirical fact is sluggish growth. Habits are not a fact, they are a possible interpretation of the fact of sluggish growth. Other interpretations might not involve habits. For example, Carroll and Slacalek (2006) show that, in aggregate data, habit formation models are indistinguishable from models without habits in which consumers have accurate knowledge of their own idiosyncratic circumstances but sticky expectations about the macroeconomy.

[3]For fuller derivations and discussion, see `Habits.pdf` in the first author's graduate macroeconomics lecture notes, available on his web page.

consumption C^* (i.e. incorporating any purely transitory elements) can be written as

$$\Delta \log C^*_{t+1} = \beta_0 + \beta_1 \Delta \log C^*_t + \zeta_{t+1}. \quad (2)$$

Sommer (2002) points out that if the transitory element of consumption spending is confined to a quarter (which he argues is a reasonable assumption both for measurement error and for hurricanes), while the process for true consumption follows (1) (possibly with $\chi = 0$, which corresponds to the random walk model), the error process for measured consumption will be an MA(1) with a negative coefficient. In these circumstances direct estimation of (2) on NIPA data will yield an estimate of β_1 that is a downward-biased estimate of the habit parameter χ, where the size of the bias will depend on the magnitude of the transitory component of expenditures.

Sommer shows that instrumental variables estimation using instruments dated $t - 2$ or earlier should largely overcome these problems. In addition, he presents IV estimates that 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 alone).

This is a case where the first stage of the two-stage least squares instrumenting process is as interesting as the second stage. The first stage, in principle, is just a regression of the form

$$\Delta \log C_t = Z'_{t-1}\eta + \nu_t,$$

where Z_{t-1} is the set of instruments. There is an extensive literature, provoked originally by Hall (1978), that seeks variables that are good at predicting consumption growth. In Hall's original paper, changes in stock prices were identified as one of the few variables with predictive power; subsequently a variety of other useful variables were identified, including various measures of interest rates, consumer sentiment, and lagged income growth.

2.1 Estimating the Wealth Effect

To estimate the wealth effect, we have to modify the Sommer methodology in several directions. First, the ultimate goal here is to obtain an estimate of the marginal propensity to consume out of wealth. But (1) 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 stock and a nonstock component. If the null hypothesis is equal MPCs out of the two components then the coefficients on their log changes will *not* be identical unless stock and nonstock wealth are the same size in every period.

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.[4] That is, if we define

$$\begin{aligned} \partial C_t &= (C_t - C_{t-1})/C_{t-5} \\ \partial W_{t-1} &= (W_{t-1} - W_{t-2})/C_{t-5} \end{aligned}$$

and so on,[5] then a first-stage regression of the form

$$\partial C_t = \alpha_0 + \alpha_1 \partial W_{t-1} \tag{3}$$

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 W^S and W^N are the stock and nonstock components of wealth, a first-stage regression of the form

$$\partial C_t = \alpha_0 + \alpha_1 \partial W^S_{t-1} + \alpha_2 \partial W^N_{t-1} \tag{4}$$

yields directly comparable estimates of relative MPCs.

Regressions of the form (3) or (4) pass all the standard tests of instrument validity and therefore justify estimation of an IV equation of the form

$$\partial C_{t+1} = c_0 + \chi \partial C_t + \xi_{t+1}. \tag{5}$$

where c_0 is an unimportant constant.

Given an initial (current-quarter) MPC out of wealth of μ and a serial correlation coefficient χ for ∂C, the usual infinite horizon formula implies that the ultimate effect on the level of consumption (the "long-run MPC") from a unit innovation to wealth is

$$\text{MPC}^{LR} = \frac{\mu}{1-\chi}.$$

As a digression, this seems a good place to point out that "the long-run MPC" is a concept of questionable fundamental validity, because in the long-run the amount of wealth is endogenous with respect to consumption choices; indeed, one interpretation

[4] Because 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

[5] We are using the ∂ symbol for one-quarter changes in wealth because later we will be defining ∂W as a weighted average of a year's worth of wealth changes.

of the cointegration discussion above is that the only sensible definition of the "long-run MPC" is that it is zero.

Our interpretation of the econometric object we call the "long-run 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. Thus the distinction between what we are presenting as a long-run 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 "long-run MPC," would have been to directly report the relevant coefficient estimates on one-quarter-lagged ∂W from the first-stage regressions. If that MPC had been α then the fact that we should have $\alpha = \chi\mu$ implies that the long-run MPC could have been estimated from

$$\mathrm{MPC}^{LR} = \frac{\alpha}{\chi(1-\chi)},$$

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.

Unfortunately, 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 reporting them directly.[6] 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:

$$\Delta C_{t+1} \approx \mu\chi(\Delta W_t + \chi\Delta W_{t-1} + \chi^2\Delta W_{t-2} + \chi^3\Delta W_{t-3}) + \eta_{t+1}.$$

Now define

$$\partial W_t = (\Delta W_t + \chi\Delta W_{t-1} + \chi^2\Delta W_{t-2} + \chi^3\Delta W_{t-3})/C_{t-4}$$

and since similarly $\partial C_{t+1} = (C_{t+1} - C_t)/C_{t-4}$ this leads to an approximate equation for ∂C and ∂W of the form

$$\partial C_t = c_0 + \alpha_1 \partial W_{t-1}. \quad (6)$$

[6] The estimates are not enormously sensitive—they typically imply long-run MPCs between 0.02 and 0.1.

Under the assumption that the dynamic model of consumption is right, the coefficient estimate on ∂W_t should be the immediate (first-quarter) MPC out of an innovation to wealth.

Thus, the estimate of the long-run MPC out of wealth reported in table 2 is given by

$$\mathrm{MPC}_n^{LR} \quad = \quad \frac{\alpha_n}{\chi(1-\chi)}. \tag{7}$$

for the α_n corresponding to the respective measure of wealth.

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

1. Estimate (5) by IV, generating the estimate of χ reported in table 2.

2. Construct the estimate of ∂W as per (6).

3. Estimate (6) or the corresponding equation for the other instrument sets, yielding the estimate of the short-run MPC contained in table 1.

4. Construct the estimate of the long-run MPC for table 2 via (7).

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 long-run MPC out of wealth shocks. However, when only 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. On the other hand, when a few lags in the equation are used the sum of the coefficients on the lags tends to yield similar short-run coefficients, but is harder to summarize. The result was the compromise represented by table 1.

2.2 Estimation Results

As a baseline, the first row of table 1 presents the estimation results of the regression (6) of the change in consumption ∂C_{t+1} on a weighted average of the change in wealth over the prior year ∂W_t. 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, 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 R^2 from the regression of ∂C_t on a constant and $\partial W_{t-1}, \ldots, \partial W_{t-4}$ not reported in table 1).[7]

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 adding 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. The latter 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.

The third row regresses the consumption change on the change in stock and nonstock wealth separately; the point estimate of the effect of nonstock wealth is more than twice as large as the coefficient on stock wealth (which is close to the original estimate of the effect of total wealth). However, the coefficient on nonstock wealth is much less precisely estimated than the coefficient on stock wealth, and a statistical test indicates that the hypothesis that the two coefficients are actually equal cannot be statistically rejected at the 95 percent significance level (this is what "(Accepted)" means in the second-to-last column). One reason the coefficient on nonstock wealth is harder to pin down is that nonstock wealth varies considerably less than stock wealth, as shown in figure 1.

The final row presents the preferred specification, in which stock and nonstock wealth effects are examined separately from the other explanatory variables. Results are broadly what would be expected: Both coefficients are substantially smaller, and the coefficient on nonstock wealth is about twice as large as that on stock wealth, but the difference between the two coefficients is not statistically significant. However, the coefficient on nonstock wealth is statistically significantly different from zero, at the 0.077 percent level.

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

The results in this table are not the bottom line, because they reflect only the next-quarter effect on consumption growth. To obtain the long-run MPCs, we need to estimate equation (5) and apply formula (7). Results of these calculations are reported in table 2.[8]

The first column shows that all models find a very substantial, and highly statistically significant, amount of momentum 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 from the authors), 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 ∂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.[9]

The last two columns report the estimated "long-run" MPCs out of stock and nonstock wealth. When the MPCs are permitted to differ for stock and nonstock wealth, the higher short-run MPCs out of nonstock wealth from table 1 translate into higher long-run MPCs here, with the preferred model estimate (the last row) of a long-run MPC out of nonstock wealth of 9 cents on the dollar.

One intuition for why the MPC out of stock wealth is substantially lower than that out of nonstock wealth is shown in Figure 1. Stock wealth is considerably more volatile than nonstock 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 consumption fluctuations, so the coefficient is not estimated to be very large. [10]

As usual, theory does not give a clear answer about whether we should expect a larger effect for stocks or houses. It might seem that more liquid assets (such as

[8] It may be surprising that current income growth is not included as an additional regressor, *a la* Campbell and Mankiw (1989). But Sommer (2002) reports that when lagged consumption growth and current income growth, instrumented by the same variables, are included together, the coefficient on lagged consumption growth remains statistically unchanged (a bit lower, but still greater than 0.5), and current income growth is not statistically significant, with a point estimate that is much lower than in the usual Campbell–Mankiw regressions. The Sommer results can be confirmed using the model presented here.

[9] To address a final concern: Using 2003 vintage data, as opposed to Hall's 1976 vintage, there are plenty of instruments with a lot of predictive power—the Fed funds rate by itself generates an $\bar{R}^2$ greater than 0.16.

[10] Figure 1 actually shows levels of wealth rather than differences ∂W used in the regressions. However, the volatility in levels is transferred into differences. Consequently, it turns out that the standard deviation of the stock wealth measure ∂W^S is about three times as large as that of ∂W^N.

stocks) are more likely to be used to smooth consumption. Consequently, consumption should respond more strongly to shocks to liquid assets. On the other hand, the distribution of wealth across households also affects the MPC estimated in aggregate data. The median dollar of stock wealth is held by a considerably wealthier household than the median dollar of housing wealth. Economic theory, empirical evidence, and common sense support the proposition that the marginal propensity to consume out of wealth is smaller for richer households. So even if the MPC out of stock wealth is greater than that out of housing wealth for any individual household, in the aggregate an increase in housing wealth could produce a larger boost to consumption than an increase in stock wealth because the housing wealth increase is concentrated among households with a higher overall MPC. In addition, it is plausible to imagine that consumers think of house price increases as a more reliable or permanent increase in wealth than an equivalent dollar change in stock market values, whether or not such a view is supported by econometric evidence or economic theory. Such a set of beliefs could justify a substantially larger effect of housing wealth than stock wealth on consumption.

Our preferred estimates pin down the "long-run" MPC out nonstock wealth at 0.09, more than twice as high as the MPC out of stock wealth of 0.04. This suggests that the theoretical arguments for the larger size of housing wealth outweigh the ones in the opposite direction.

2.3 Comparison with the Existing Empirical Work

The work most closely related to ours is Case et al. (2005), which provides estimates from both a panel of developed countries (since 1975) and a panel of states within the US. Using annual data, Case et al. (2005) find a highly statistically significant estimate of the MPC out of housing wealth in the US of around 0.03–0.04. In contrast, the Case et al. 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 stock wealth.

Unfortunately, the econometric evidence does not speak with one voice. An IMF study by Ludwig and Sløk (2002) estimates a larger effect of stock 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 stock wealth effect is stronger (and in some neither was significant). And a study by Dvornak and Kohler (2003) modelled closely on the Case et al. study but using Australian state-level data finds a larger stock wealth effect than housing wealth effect.

There are reasons to be skeptical of results based on macroeconomic or regional data. Perhaps the foremost is 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. Thus, to isolate a "pure" housing wealth effect, one would want data on spending by individual households before and after some truly exogenous changes in their house values, caused for example by the unexpected discovery of neighborhood sources of pollution.

Recent studies by Disney et al. (2003) and Campbell and Cocco (2006) seem to represent the closest approximations to such an ideal microeconomic dataset. Disney et al. use data on spending patterns for a set of British households, along with county-level indicators of house prices. They find a median marginal propensity to consume out of housing wealth in the range 0.09–0.14 during the recent house price boom in Britain.

Campbell and Cocco (2006) use British data from the UK Family Expenditure Survey and the Nationwide data set of regional home price indexes to investigate the wealth effect for various demographic groups. 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. Additional microeconometric estimates of the MPC out of wealth are reported in Engelhardt (1996), Juster et al. (2001), Lehnert (2001) and Levin (2001).

On balance, most macro- and microeconometric studies provide evidence that the, medium-run MPC (after 3 years) out of housing wealth in the range of 0.04–0.10, comparable to our estimates from aggregate US data.

3 Comparison of Our Methodology with the Cointegration Approach

Our approach described in section 2 provides an alternative way to estimate the wealth effect without having to appeal to the strong assumptions which are embedded in the currently popular cointegration methodology.[11] The cointegration methodology consists of estimating the cointegrating regression between (logarithms of) consumption, labor income and wealth. This work typically invokes as a theoretical justification the following log-linear approximation to the consumer's intertemporal

[11] The number of papers applying the cointegration methodology has recently risen considerably. References include Bertaut (2002), Byrne and Davis (2003), Fernandez-Corugedo et al. (2003), Pichette and Tremblay (2003), Catte et al. (2004), Lettau and Ludvigson (2004) and Hamburg et al. (2005).

budget constraint (based on Campbell and Mankiw, 1989)

$$\log C_t - \frac{\Gamma}{\Psi}\log Y_t - \left(1-\frac{\Gamma}{\Psi}\right)\log W_t = \sum_{i=1}^{\infty}\gamma^i\left(\frac{\Gamma}{\Psi}\Delta\log Y_{t+i} - \Delta\log C_{t+i} + \left(1-\frac{\Gamma}{\Psi}\right)r_{t+i}\right), \tag{8}$$

where C_t is consumption, Y_t labor income, W_t household asset holdings (wealth), $1-\Gamma/\Psi$ is the steady-state ratio of savings invested in physical assets and $\gamma = 1-\exp(\overline{\log(C/W)})$ with $\overline{\log(C/W)}$ being the log of the steady-state consumption–wealth ratio.

Equation (8) is often presented as a preference-free relationship implied solely by the budget constraint, embodying minimal theoretical restrictions, e.g. Lettau and Ludvigson (2004), p. 280:

> Importantly, however, the empirical approach described next does not require imposing such additional [preference] structure. The empirical results we obtain exploit only cointegration, a phenomenon that can be motivated by the logic of a simple budget constraint identity, applicable to a wide variety of theoretical structures.

As hinted above, the problem with this argument is that the underlying theory does not suggest a stable γ unless all aspects of the economy are perpetually unchanging. The assumption is effectively that the target ratio of market wealth to total wealth Γ/Ψ, or equivalently the ratio of market wealth to human wealth is stationary. Unfortunately for the cointegration approach, neither partial nor general equilibrium theory provides a plausible rationale for this stationarity assumption. For example, in the classic Hall (1978) model with a quadratic utility function the level of consumption follows a random walk. Since consumption is proportional to total wealth, the level of total wealth must also follow a random walk. But if income (and therefore human wealth) trends upward, it must be the case that Γ/Ψ is falling. So even the benchmark Hall model does not imply stability in the required sense.

More deeply, in any sensible macroeconomic model, the long-run relationship between (the logs of) consumption, income and wealth is bound to be determined by the steady state interest rate, productivity growth rate, population growth rate, tax rates, insurance arrangements, and other parameters. If any of these parameters change permanently over the time frame examined, any efforts to find a unique, stable cointegrating relationship between the three variables is misguided.[12] This is of course particularly problematic if one hopes to separate a long-run cointegrating

[12] Empirically, there is strong evidence of persistent changes in the mean of productivity growth in the US. The average labor productivity growth in the non-farm business sector was 2.6 percent between 1955 and 1972, 1.5 percent between 1973 and 1994 and 2.7 percent between 1995 and 2004. For additional evidence see e.g. Hansen (2001).

relationship from persistent transitional dynamics toward the steady state. In that case to make reliable estimation and inference it is crucial to have long spans of structurally stable data, which are unlikely to be available in practice.

The empirical evidence on the existence of unique, stable cointegrating relationship between consumption, income, and wealth is rather mixed. Some of the papers cited above argue that there is no evidence against the stability of the cointegrating relationship, but others, including Hahn and Lee (2001), Rudd and Whelan (2002) and Slacalek (2004) argue that with properly constructed data the relationship does not exist or is structurally unstable.

In addition, the cointegration approach effectively postulates that consumption growth depends *only* on its own lags and the lags of income and wealth. The error-correction counterpart of cointegrating regression (8) is

$$\Delta \log C_t = \nu_c + \alpha_c \hat{ec}_{t-1} + \Gamma_c'(L)\big(\Delta \log C_{t-1}, \Delta \log Y_{t-1}, \Delta \log W_{t-2}\big)' + e_{c,t}, \qquad (9)$$

where $\hat{ec}_{t-1}$ is the deviation from the long-run trend (cointegrating residual). It is, however, well-known that additional variables, such as consumer sentiment or measures of liquidity constraints predict consumption growth (see e.g. Carroll et al., 1994 and Bacchetta and Gerlach, 1997). As a result, estimates based on equations like (9) potentially suffer from omitted variable bias.

4 Conclusion

Housing price fluctuations apparently have substantial effects on consumer spending. The immediate (first-quarter) impact is likely to be 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 it probably accumulates to the 4–10 cent range. These figures are consistent with evidence from micro data and the experience across US states. Whether the housing wealth effect is substantially larger than the stock wealth effect is more uncertain; while the bulk of the evidence seems to point in that direction, the estimated size of the differences is not large enough (in US aggregate data) to yield confidence in the conclusion. For monetary policy purposes, these results suggest that it is important 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 US consumption and residential investment spending in the early 2000s even as the stock market suffered a historic decline.

More importantly, the risks of the opposite experience are also worth noting. While in most places there seem to be good fundamental reasons for the rise in housing wealth over the last few years, the housing price dynamics in some areas might be driven by bubble components. If these components grow large and decline abruptly, our results suggest that consumption will be affected substantially.

References

Bacchetta, Philippe, and Stefan Gerlach (1997), "Consumption and Credit Constraints: International Evidence," *Journal of Monetary Economics*, 40, 207–238.

Bertaut, Carol (2002), "Equity Prices, Household Wealth, and Consumption Growth in Foreign Industrial Countries: Wealth Effects in the 1990s," IFDP working paper 724, Federal Reserve Board.

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.

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 Olivier J. Blanchard and Stanley Fischer, editors, *NBER Macroeconomics Annual*, 185–216, MIT Press, Cambridge, MA.

Carroll, Christopher D. (2004), "Housing Wealth and Consumption Expenditure," *Paper Prepared for Academic Consultants Meeting of Federal Reserve Board, January 2004*, available at `http://econ.jhu.edu/people/ccarroll/papers/FedHouseWealthv2.pdf`.

Carroll, Christopher D., Jeffrey C. Fuhrer, and David W. Wilcox (1994), "Does Consumer Sentiment Forecast Household Spending? If So, Why?" *American Economic Review*, 84(5), 1397–1408.

Carroll, Christopher D., and Jirka Slacalek (2006), "Sticky Expectations and Consumption Dynamics," mimeo, Johns Hopkins University.

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.

Catte, Pietro, Nathalie Girouard, Robert Price, and Christophe Andre (2004), "Housing Market, Wealth and the Business Cycle," working paper 17, OECD.

Disney, Richard, Andrew Henley, and Jonathan Heathcote (2003), "House Price Shocks, Negative Equity and Household Consumption in the UK," mimeo, 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 Discussion Paper 23.

Dynan, Karen E. (2000), "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 90(3), 391–406.

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), "Habit Formation in Consumption and Its Implication for Monetary-Policy Models," *American Economic Review*, 90(3), 367–390.

Girouard, Nathalie, and Sveinbjörn Blöndal (2001), "House Price and Economic Activity," OECD Economics Department Working Paper 279.

Hahn, Jaehoon, and Hangyong Lee (2001), "On the Estimation of the Consumption-Wealth Ratio: Cointegrating Parameter Instability and its Implications for Stock Return Forecasting," mimeo, Columbia University.

Hall, Robert E. (1978), "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis," *Journal of Political Economy*, 86(6), 971–987.

Hamburg, Britta, Mathias Hoffmann, and Joachim Keller (2005), "Consumption, Wealth and Business Cycles in Germany," discussion paper 16, Deutsche Bundesbank.

Hansen, Bruce E. (2001), "The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity," *Journal of Economic Perspectives*, 15(4), 117–28.

Juster, F. Thomas, Joseph Lupton, James P. Smith, and Frank Stafford (2001), "The Decline in Household Saving and Wealth Effect," mimeo, University of Michigan.

Lehnert, Andreas (2001), "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 (2001), "Are Assets Fungible? Testing the Behavioral Theory of Life Cycle Saving," *Journal of Economic Behavior and Organization*, 36, 59–83.

Ludwig, Alexander, and Torsten Sløk (2002), "The Impact of Changes in Stock Prices and House Prices on Consumption in OECD Countries," IMF Working Paper 1.

Muellbauer, John (1988), "Habits, Rationality and Myopia in the Life Cycle Consumption Function," *Annales d'Economie et de Statistique*, 9, 47–70.

Pichette, Lise, and Dominique Tremblay (2003), "Are Wealth Effects Important for Canada?" working paper 30, Bank of Canada.

Reis, Ricardo (2004), "Inattentive Consumers," mimeo, Princeton University.

Rudd, Jeremy, and Karl Whelan (2002), "A Note on the Cointegration of Consumption, Income, and Wealth," FEDS working paper 38, Federal Reserve Board.

Slacalek, Jirka (2004), "International Evidence on Cointegration between Consumption, Income, and Wealth," mimeo, Johns Hopkins University.

Sommer, Martin (2002), "Habits, Sentiment and Predictable Income in the Dynamics of Aggregate Consumption," working paper number 458; updated 2006, Johns Hopkins University, available at `http://econ.jhu.edu/pdf/papers/wp458_version2006.pdf`.

Appendix: Description of Data

Consumption Total personal consumption expenditure; 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.

Stock wealth 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.

Nonstock wealth Net worth − Stock 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, 1960Q1–2004Q3.

The archive with programs and data used in this paper can be downloaded at `http://www.econ.jhu.edu/people/ccarroll/`.

Figure 1: Components of Household Wealth

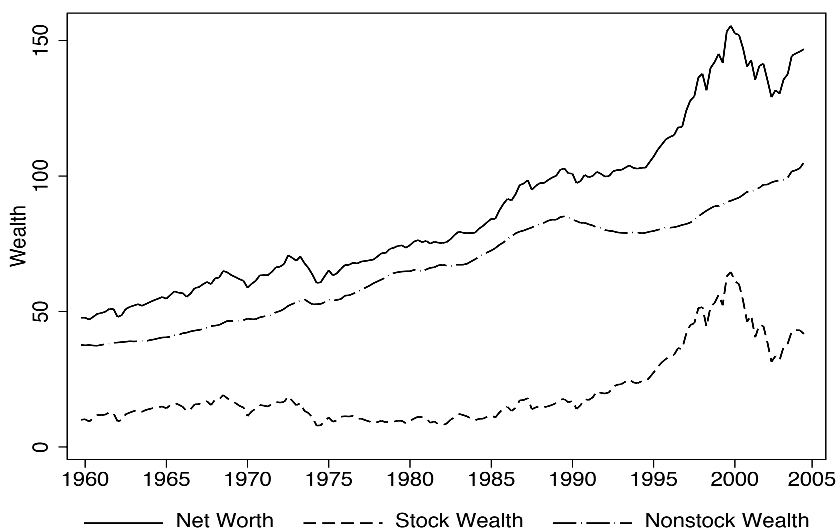

Note: Per capita real wealth figures in thousands of year 2000 dollars. Net worth is our measure of total wealth.

Table 1: Short-Run Effect of Wealth on Consumption

$$dC_t = \alpha_0 + \alpha_1 dW_{t-1} + \alpha_2 dW^S_{t-1} + \alpha_3 dW^N_{t-1} + \alpha_4 MU_{t-1} + \alpha_5 FF_{t-1}$$

Next-Quarter Effect of \$1 Change in Wealth			Extra Variables			
Total dW_{t-1}	Stock dW^S_{t-1}	Nonstock dW^N_{t-1}	Unemp Exp MU_{t-1}	Fed Fund FF_{t-1}	Test of $dW^S = dW^N$	$\bar{R}^2$
0.0168*** (0.0044)						0.128
0.0107*** (0.0032)			0.081** (0.033)	−0.457** (0.208)		0.232
	0.0157*** (0.0046)	0.0381*** (0.0113)			0.078	0.136
	0.0081*** (0.0026)	0.0182** (0.0083)	0.073** (0.036)	−0.492** (0.214)	0.239	0.230

Notes: Sample period is 1960Q1–2004Q3. 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 0.0167 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 2: Consumption Growth Momentum and the Long Run MPC

$$dC_{t+1} = c_0 + \chi \mathbf{E}_{t-1} dC_t + \zeta_{t+1}$$

Variables used to forecast $\mathbf{E}_{t-1}dC_t$	Consumption Growth Momentum Coefficient χ	Implied Long-Run MPC out of		
		Total W	Stock W^S	Nonstock W^N
W	0.59^{**}	0.069		
	(0.23)			
W,	0.78^{***}	0.061		
MU, FF	(0.14)			
W^S, W^N	0.47^{**}		0.063	0.153
	(0.20)			
W^S, W^N,	0.73^{***}		0.041	0.091
MU, FF	(0.13)			

Notes: Sample period is 1960Q1–2004Q3. Standard errors are in parentheses. {*,**,***}=Statistical significance at {10,5,1} percent. The long-run MPCs are calculated from the formula $\alpha_n/\chi(1-\chi)$ where α_n is the corresponding next-quarter MPC estimated in table 1. Standard errors for all equations are heteroskedasticity and serial-correlation robust. When more instruments are used to forecast dC_t (for example, the Fed funds rate and the change in unemployment over the previous year), the estimate of ρ tends to rise further and the standard error falls further. The measure of the change in wealth used for the regressions is the ∂W measure defined in the text, as this can be measured without an estimate of ρ, unlike the dW measures used in the previous table.