Housing Wealth and Consumption Expenditure

Christopher D. Carroll*
ccarroll@jhu.edu

January 30, 2004

Abstract

This memo considers the effect of housing wealth fluctuations on consumption expenditures. While empirical evidence from macroeconomic, regional, and microeconomic data varies, on balance most of the evidence seems consistent with a medium-run (after 3 years or so) housing wealth MPC in the neighborhood of 5-10 cents. The paper presents its own methodology for estimating the housing MPC, and concludes that the immediate (next-quarter) MPC from a change in housing wealth is around 1.5 cents, with a final long-run effect of about 9 cents. Finally, the paper estimates that the growth in housing wealth since 2000q1 left the level of aggregate consumption about 2.2 percentage points higher by 2003q3 than it would have been if housing wealth had been flat over the intervening period, possibly providing a partial explanation for the surprising strength of spending in the wake of the stock market declines over the last few years.

Paper prepared for Academic Consultants' meeting of the Board of Governors of the Federal Reserve System, January 30, 2004.
I am grateful to Emmanuel de Veirman, Sherif Khalifa, and Misuzu Otsuka for excellent research assistance in preparing this memo.

* Department of Economics, Johns Hopkins University, Baltimore MD 21218-2685.

1 Introduction

Aside from stock market prospects, there are few questions on which an economist at a dinner party is more likely to be consulted, and less likely to believed, than the direction of the housing market.

The widespread interest in house prices makes sense, because home equity makes up the bulk of wealth for the majority of households. And, on balance, the evidence from micro, regional, macro, and international sources suggests that consumers' interest in the topic is not merely academic: When house prices rise, the perceived increase in wealth seems to produce extra spending, though the exact magnitude of the housing wealth effect is hard to pin down.

Conveniently, whatever effects house price fluctuations may have on spending, those effects seem to be relatively gradual. This gives policymakers the luxury of waiting to see effects on spending; indeed, in the end, the traditional focus of monetary policy on whether aggregate demand is outstripping aggregate supply may be easier to judge than whether movements in demand are attributable to housing prices, stock prices, consumer confidence, or other factors. But at a minimum, the evidence of a substantial housing wealth effect indicates that movements in housing prices are something that policymakers should keep a close eye on.

2 Who Wins, Who Loses, Who Spends, Who Saves?

Prices of Old Master paintings depend entirely on movements in demand. New Rembrandts will not be produced no matter how high the price goes, nor will existing paintings be destroyed no matter how far their price falls. Prices of manufactured goods like drapes, on the other hand, tend to be driven to the cost of manufacture. Whether in the long run the price of the housing stock is determined more like the price of drapes or more like that of Rembrandts is a difficult question, and depends partly on particular market circumstances like the nature of zoning restrictions and the capacity of transportation networks.

This may explain the much greater volatility of housing prices in Europe than the U.S., as reflected in figure 1 (taken from Case, Quigley, and Shiller 2003, henceforth CQS); a recent survey article in *The Economist* argued that European markets present substantially greater obstacles in bringing new houses to market than prevail in most of the United States. In an international context, both the 27 percent rise in real U.S. house prices since 1995,[1] and the overall variation in housing wealth (from an index value of 80 in 1975 to about 105 in 1999 as reported in CQS) are quite moderate, and suggest that housing is more drapelike in the U.S. than elsewhere.

[1] *Economist* magazine housing markets survey, November 2003, p. 6.

Even in the U.S., though, there are individual markets where little buildable new land remains; for such places, house prices should depend on local demand. The ongoing urbanization of the U.S. population, and continuing migration toward cities like San Francisco and L.A. on the West coast, and Washington, New York, and Boston in the East, may be able to justify the rise prices in most of these areas, though bubbles cannot be ruled out in some cases. Prices in San Francisco, for example, are up 70 percent since 1995,[2] which may or may not be a sustainable reflection of Northern California's prosperity.

There are also reasons for relative house prices to rise over time even if houses are more like drapes than paintings. As famously pointed out by William Baumol, the relative price will tend to rise for goods for which productivity gains are relatively poor. Since one of the robust stylized facts of productivity research is that productivity gains in the construction industry have been meager to nonexistent (see, e.g., Slifman and Corrado 1999), economic theory would lead one to expect a rise in the relative price of housing even if zoning, transportation, and other constraints did not exist; indeed, Davis and Heathcote 2004 find that the long-term trend of house construction prices has indeed been persistently upward. (Davis and Heathcote also note that since buildings are reproducible assets, the truly Rembrandtesque component of housing wealth is the underlying land, whose value they attempt to measure separately from the value of the stock of buildings; they find that land prices have risen much faster than overall housing prices.)

Also, as with any asset, the price of housing should fluctuate with real interest rates, which have been low in the last few years. In addition, house prices may also be affected by nominal rates: The fact that most housing purchases are financed by constant nominal rate mortgages means that for a given real interest rate, a lower expected long-term inflation rate translates into increased real borrowing capacity because the profile of real mortgage payments is shifted into the future. This higher borrowing capacity could translate into higher real demand and therefore higher prices. Thus, the secular decline in inflation could be partly responsible for the increase in relative housing prices.

From the standpoint of consumer spending effects, the implications of house price fluctuations are murkier than one might suppose.

Homeowners, of course, tend to be pleased when the value of their homes rises. But if they intend to continue living in the same house indefinitely, or a similar house in another neighborhood where prices have gone up commensurately, the rise in prices does not increase their real wealth, so should not increase their spending.

Furthermore, for every current homeowner who is potentially made better off by a rise in relative home prices, there is a future homeowner (maybe currently a renter; maybe currently a child of the homeowner) who is made worse off. To the extent that parents view the house as something to leave to the kids, a rise

[2] *Economist* survey, p. 6.

in its value is not an occasion for a spending spree. And renters may need to reduce their spending if they hope ever to become owners.

Thus, judging from the effects on spending power across current and future generations, a rise in relative house prices does not increase the amount of real resources that can be spent on nonhousing goods and services in the long run. Another way to see this is to realize that a rise in house prices does not increase the economy's ability to produce nonhousing goods, and ultimately consumption of such goods must equal their production.

To say that there are no long-term effects on aggregate spending capacity, however, is not the same as saying there can be no short- or medium-term effects on actual spending. The winners from a house price increase may increase their spending today, while the losers may decrease their spending in the future.

One obvious mechanism by which a rise in housing prices can promote short-term spending is by allowing households who had borrowed close to their mortgage maximum to refinance and increase their mortgage debt in accord with the rise in their house value. Of course, if money extracted by refinancing is used for home improvements, the result is a boost to residential investment spending rather than consumption expenditure. But either kind of spending is a boost to aggregate demand in the short run.

Spending effects may also depend on whether the capital gain is realized. While sale proceeds are often plowed back into equity in a replacement home; and some evidence from Lehnert 2003, for example, shows that older sellers, who are more likely to trade down, have higher housing MPC's.

The theoretical effects of stock prices on spending are somewhat clearer. From the long-term perspective, since stocks represent assets that can be traded internationally, an increase in stock wealth is arguably a real increment to national spending power (at least, more so than an increase in housing wealth). Furthermore, stock prices should reflect expectations about future productivity growth, and thus a rise in stock wealth may be a signal of increased future prosperity, which should translate into more spending today. Also, recent theoretical work by Otsuka 2003 shows that even a modest degree of illiquidity (for example, the time and expense associated with getting an appraisal done for refinancing purposes) can substantially slow down the translation of an increase in wealth into an increase in spending. These considerations all suggest that the effects of stock prices on consumption should be larger than those of house prices, at least in the medium run.

On the other hand, 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 would be 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.

Other considerations may also matter. It is plausible to suppose 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, the effect on consumer spending depends on what consumers think, not what economists think. Such a set of beliefs could justify a substantially larger effect of housing wealth than stock wealth on consumption.

To broaden this point, how consumers behave depends on how they think about things, not on how economists think they should think about them. If Thaler 1990 is right in proposing that consumers have a set of "mental accounts" that they use to evaluate developments in different spheres of their economic lives, there is no necessary connection between the spending effects housing wealth and stock wealth.

In the end, as usual, abstract economic theorizing does not produce unambiguous conclusions. The fact that housing wealth is more broadly distributed, is viewed as more reliable than stock wealth, and can be more easily borrowed against suggests that housing wealth effects should be larger than stock wealth effects, while the fact that stock wealth arguably reflects true increases in aggregate spending power, and may be more liquid than housing wealth, suggests the reverse. One would hope, therefore, that the question could be resolved by empirical evidence.

3 Macroeconomic Evidence

The traditional approach to estimating wealth effects involves estimating regressions of the saving rate on the wealth-to-income ratio. A state-of-the-art example is the recent paper by Davis and Palumbo 2001, which estimates an overall wealth effect in the range of 0.04-0.06. Davis and Palumbo (2001) also examine whether there is any evidence of a different MPC out of stock wealth and nonstock wealth. Since home equity makes up the bulk of nonstock wealth, this can be interpreted as a test of whether there is a different MPC out of housing and stock wealth. Davis and Palumbo find a long-run MPC out of nonstock wealth of 0.08, compared to an MPC out of stock wealth of about 0.06 in the corresponding specification.[3] The difference between the two coefficients is on the edge of statistical significance.

An important drawback of this econometric approach is that it implicitly assumes that the saving rate and the real interest rate are constant over the long run. In an economy that is constantly changing, this assumption is problematic. Economic theory says that the personal saving rate should depend on the structure of the retirement system, the tax system, expected productivity

[3]Model 1b, Table 10, p. 35.

growth rates, and many other features of the economy that have changed profoundly over time. And an assumption of a constant saving rate is empirically problematic in light of either the long-term downtrend in the saving rates in most industrial countries including the U.S., or the dramatic differences in saving rates across countries (which suggest that there is no "natural" saving rate common to rich countries, and therefore perhaps no natural saving rate for any individual country either).

The usual alternative way of estimating wealth effects involves attempting to find a common trend in the logarithms of consumption, income, and total wealth. But this method also depends on an assumption of long-term stability, and in my view there are other deep statistical and conceptual problems with this approach; in addition, the method cannot be straightforwardly adapted to allow different but quantitatively comparable "MPC's" for nonstock and stock wealth.

Instead, I have developed an alternative technique based on the finding in the recent academic literature that there is a considerable degree of predictability in consumption growth over frequencies of a few quarters.

As a baseline, the first row of table 1 presents the results when the change in consumption (normalized by consumption from a year earlier) is regressed on a weighted average of the change in wealth over the prior year.[4] The regression coefficient is scaled so that it is directly interpretable as the marginal propensity to consume out of changes in wealth in the previous quarter.[5] 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. This wealth effect is highly statistically robust, but lagged wealth growth alone explains only about 13 percent of quarterly consumption growth.[6]

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, income growth rates, and a variety of other variables. However,

[4]Specifically, the wealth growth measure incorporates geometrically declining weights on lagged changes in wealth where the rate of decay is estimated as discussed below.

[5]To be exact, the regression is of $(C_t - C_{t-1})/C_{t-5}$ on a geometrically declining weighted average of $(W_{t-1} - W_{t-2})/C_{t-5}$ and its lags, where the normalization by C_{t-5}, along with the declining weights chosen for the wealth terms, ensure that the coefficient is roughly interpretable as an MPC.

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

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, 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 more than 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 about 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 (this is what "(Accepted)" means in the second-to-last column).

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 0.053 percent level.

The results in this table are not the bottom line, because they reflect only the next-quarter effect on consumption growth. Recent work by Fuhrer 2000, Sommer 2001, Gabaix and Laibson 2001 and others has found that consumption reacts sluggishly to shocks, so we should not expect the immediate adjustment reported in table 1 to correspond to the final effect.

The canonical model of sluggish adjustment is an equation of the form

$$dC_{t+1} = c_0 + \rho dC_t + \epsilon_{t+1}, \tag{1}$$

where ρ is the serial correlation coefficient ("momentum") and c_0 is an unimportant constant. When equations of this form are estimated directly on U.S. NIPA consumption data, ρ is usually not particularly large (though it is statistically significant). However, Sommer 2001 points out that there are good reasons to expect transitory fluctuations in measured consumption expenditures. Aside from measurement error (which is considerable at the quarterly frequency), transitory events such as bad weather or natural disasters like the recent earthquake in

California add high-frequency noise to lagged consumption that econometrically obscures the underlying momentum inherent in consumption growth. Sommer proposes a simple instrumental variables estimation technique to overcome this problem, which essentially involves estimating equations like those summarized in table 1 and substituting predicted consumption growth for the (noisy) actual consumption growth.

Results of this experiment for each of the models estimated in table 1 are presented in table 2, along with the implied "long-run" MPC's out of the two measures of wealth (in the last column).[7]

The first column shows that all of the 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. (The result of a highly statistically significant serial correlation coefficient in the 0.6-0.9 range also holds when the measure of consumption expenditures is restricted to spending only on nondurable goods).

The last two columns report the estimated "long-run" MPC's out of stock and nonstock wealth. (Note that the reason "long-run" is in quotes is that these results cannot be taken seriously as estimates over very long periods (say, more than ten years). They should really be interpreted as summarizing the effects that should be expected within about three years). When the MPC's are permitted to differ for stock and nonstock wealth, the higher short-run MPC's out of nonstock wealth from table 1 translate into higher long-run MPC's here, with the preferred model estimate (the last row) of a long-run MPC out of nonstock wealth of almost 9 cents on the dollar.

3.1 Other Evidence

For purposes of this memo, the foregoing results are disappointing because the MPC out of nonstock wealth is not statistically different from that out of stock wealth.

This could be because the comparatively modest size of movements in aggregate U.S. house prices provides insufficient variation to estimate a house price effect precisely. As noted earlier, there is much more variation in house prices in some other countries, or within regions inside the U.S., than there is in the aggregate over time for the U.S. as a whole. Several studies examine these alternative data in an attempt to obtain more statistically reliable measures of the effects

[7] It may be surprising that current income growth is not included as an additional regressor, *a la* Campbell and Mankiw 1989. But Sommer 2001 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 in this memo.

of housing and stock prices. The best of these is the previously-cited study by Case, Quigley, and Shiller 2003, which provides estimates from both a panel of developed countries (since 1975) and a panel of states within the U.S.

The CQS results that are most comparable to those presented here (their table 3) are estimated using data from U.S. states (right panel of the table). Using annual data, they find a highly statistically significant estimate of the MPC out of housing wealth of around 0.03-0.04. This is the first-year effect, which is appropriately intermediate between the first-quarter MPC's estimated in table 1 and the long-run MPCs calculated in table 2. However, in contrast with the results presented above, the CQS estimate of the MPC out of stock market wealth is small and statistically insignificant, and the coefficient on housing wealth is estimated to be highly statistically significantly larger than the coefficient on stock wealth. [8]

Results from the panel of countries are qualitatively similar: a large and highly statistically significant housing wealth effect, but no stock wealth effect. However, the coefficient estimates of the housing wealth effect are about four times as large as in the U.S. state data. This is puzzling; given that it is easier to borrow against housing wealth in the U.S., one might expect a higher coefficient in the U.S.

CQS also estimate housing and stock wealth effects in a variety of other ways, and consistently find a much larger housing wealth effect.

Unfortunately, as usual the econometric evidence does not speak with one voice.

An IMF study by Ludwig and Sløk 2004 found a larger effect of stock wealth than housing wealth in a panel of 16 OECD countries, and also found some evidence of an increase in wealth effects over time. Girouard and Blöndal 2001 of the OECD failed to find consistent results across countries: In some, the housing wealth effect was larger, while in others the stock wealth effect was larger (and in some neither was significant). And a study by Dvornak and Kohler 2003 modeled closely on the CQS study but using Australian state-level data found a larger stock wealth effect than housing wealth effect. (Though their point estimate of the housing wealth effect was about 3 cents, not far from the CQS figure cited above).

4 Microeconomic Evidence

There are many reasons to be skeptical of results based on macroeconomic or regional data. Perhaps the foremost is that movements in asset prices are not

[8]The inability of CQS to find a stock wealth effect may reflect problems in their stock wealth data. Because no direct data on holdings of stock wealth by state exist, CQS rely on data on mutual fund ownership, which they plausibly argue is the best available data. But it is possible that the change in the geographical variation in mutual fund ownership does not entirely capture the relevant dynamics of regional stock portfolios.

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. These points apply even to state-level analysis. Indeed, Greene 2002 shows that fluctuations in the prices of companies headquartered in Northern California produce measurable effects on the prices of homes there, while Jud and Winkler 2002 find city-level house prices are strongly related to population growth, income growth, and other city-level variables.

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 (or unexpected cleanup) of neighborhood sources of pollution.

A recent study by Disney et. al. 2002 seems to represent the closest approximation to such an ideal microeconomic dataset. This study uses data on spending patterns for a set of British households, along with county-level indicators of house prices. While even at the county level there are likely to be problematic "macroeconomic" effects, at the very least it seems fair to say that such effects should be less than for national data.

Disney et. al. 2002 find a median marginal propensity to consume out of housing wealth in the range 0.01-0.03 during the house price boom of the 1990s in Britain. Interestingly, they find a much stronger effect of house prices on consumption for households who start out in a position of "negative equity": That is, people who had bought a house whose value had subsequently declined to be less than the mortgage debt owed. For such people, a revival in house prices seems to bring a substantial surge in spending (the median MPC estimates are in the range 0.04-0.06), presumably reflecting a revival in their spending from depressed levels when house prices fell.

These results jibe with Engelhardt's 1996 estimate of an MPC of 0.03 for the U.S. over the period 1984-1989, and with more recent estimates by Juster, Lupton, Smith, and Stafford 2001, although the latter paper finds a much larger MPC (0.17) out of stock price changes, in contrast with Levin 1998, who finds a lower MPC out of illiquid assets than of liquid ones. Lehnert 2003 finds an MPC of about 0.04-0.05.

On the other hand, using a combination of city-level and micro data, Hoynes and McFadden 1997 found that households who had experienced housing capital gains actually increased their saving rather than their spending (which could reflect the fact that, unlike the other micro studies, Hoynes and McFadden's sample included renters).

On balance, though, the microeconometric evidence does not provide much evidence against the proposition that there is a medium-run MPC (after 3 years) out of housing wealth in the range of 0.04-0.10. It does suggest skepticism about very high estimates of more than about 0.10, at least if such high MPCs are supposed to apply within a period of a year or two after the change in housing wealth.

5 Other Issues

For the purposes of monetary policy analysis, there are a few other issues that do not fit neatly into the foregoing discussion but should be addressed because of their current relevance.

A particularly interesting question is what the simple model in section 3 has to say about the path of consumption expenditures in the period since the stock market peak. The fact that housing wealth continued to rise in recent years even as stock wealth was falling suggests the potential that the model could help explain the surprising strength of consumption spending during and after the recent recession. To investigate this question, a simulation was done of an alternative history, starting in 2000q1, in which stock market wealth was assumed to have repeated its true historical path, but the change in housing wealth was set to zero. The model implies that by 2003q3 the level of consumption expenditures was roughly 2.2 percentage points higher than it would have been had housing wealth remained flat.[9]

Another question is whether consumption has been boosted by the recent waves of mortgage refinancing motivated by the drop in long-term nominal interest rates. Intuitively, refinancers who keep the same mortgage duration will increase their monthly cash flow by an amount related to the difference in interest rates between the old and new mortgage. And some refinancers opt to increase the amount of their debt, freeing up substantial amounts of resources that can be used for current spending. On the other hand, it is often forgotten that for every winner in the refinancing game there is a loser: The holder of the high-interest mortgage that is retired experiences a capital loss of equivalent size to the gain for the refinancer. While it is likely that the MPC's of the losers are smaller than those of the winners (at least in the short run), the MPC's of the losers are probably not zero.

Even leaving aside the behavior of the losers, judging by the experience of previous waves of refinancing, the dramatic refinancing activity last year is likely to have had only a modest impact on consumption expenditures. Freddie Mac reports[10] that only about a third of consumers refinancing mortgages in 2003 extracted any extra "cash out" from the refinancing (compared to historical figures typically in excess of 50 percent), and the median age of the refinanced loan was less than two years. Thus, ecent refinancers are likely to have been

[9]It must be admitted that the difference between the coefficients on housing wealth and stock wealth contributes little to the result of a higher level of consumption; the higher consumption is attributable mainly to the higher path of overall wealth reflected in historical experience compared to what would have happened to wealth if housing wealth had remained flat.

[10]http://www.freddiemac.com/news/finance/refi_archives.htm

mostly the kind of financially sophisticated consumers who are likely to have low marginal propensities to consume. Careful calculations by Canner, Dynan, and Passmore 2002 make a persuasive case that the effects of refinancing in 2001 and the first half of 2002 were probably boosted consumption growth by at most 0.25 percent over the year ending in March 2002. Similar calculations would probably produce a comparably small estimate of the boost to consumption growth in 2003. Indeed, the effect on residential investment expenditures was probably substantially larger than the effect on PCE, but this can be interpreted as part of the usual channel by which stimulative monetary policy affects investment spending. In sum, the effect on consumption of rising housing prices, which affect all homeowners, has probably been much larger than the effect of the refinancing boom, which affected only the subset of homeowners who refinanced.

6 Conclusion

Taken as a whole, there seems to be enough evidence from enough different sources to conclude that over frequencies of a year or more, there is an economically important effect of housing price fluctuations on consumer spending. The immediate (first-quarter) effect is likely to be small (in the preferred model above, the immediate quarterly MPC was estimated to be about 1.5 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 U.S. states. Whether the housing wealth effect is substantially larger than the stock wealth effect is more uncertain; while the evidence seems to point in that direction, the estimated size of the differences is not large enough (in U.S. aggregate data) to yield confidence in the conclusion. Micro data yields a mixed picture on this question, though the results of Case, Quigley, and Shiller (2003) do point strongly to a larger housing than stock wealth effect.

For monetary policy purposes, these results suggest that it is likely to be 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 consumption and residential investment spending over the past three years even as the stock market suffered a historic decline.

In closing, the risks of the opposite experience are also worth noting. While there seem to be good fundamental reasons for the rise in U.S. house prices over the last few years (at least in most places), bubbles are never easy to perceive in real time. If, for example, a steep rise in interest rates led to a sharp decline in house prices, the balance of the risks to consumption might shift substantially more in the negative direction than might be indicated solely by the contemporaneous movements of stock prices.

Appendix

6.1 What's Wrong with Cointegration in this Context?

Cointegration analysis is the standard method of estimating the short- and long-run marginal propensities to consume out of wealth, yet this memo does not report the results from any cointegrating regressions. This reflects my conviction that cointegration analysis is deeply problematic if interpreted as a method for extracting a "long run MPC." There are two, closely related, problems: First, in order for any of its conclusions to be valid, a cointegration analysis requires stability of the cointegrating vector. And, second, consistent estimation requires exogenous variation in the independent variables. Yet the level of wealth is intimately connected with the interest rate, which in turn should be related to the marginal propensity to consume.

In the context of the usual wealth effect regressions, the stability requirement translates into an assumption that there be a "natural" personal saving rate that the economy always reverts to. This is problematic, since (as noted in the text), the saving rate should depend on many features of the economy that have clearly changed profoundly in the U.S. over the relatively short period for which the cointegration analysis is performed. Furthermore, a recent literature, with contributions by Hahn and Lee 2001, Brennan and Xia 2002, and others has found considerable evidence of instability in such cointegrating vectors.

As one example of the pitfalls of cointegration analysis, consider a standard Ramsey/Cass-Koopmans economy with a Cobb-Douglas aggregate production function. For simplicity, assume there is no underlying technological process. It can be shown (derivations available from the author) that, defining lower case variables as the upper case version divided by labor income, in long-run steady-state the relationship between c and k will be captured by

$$c = 1 + rk. \tag{2}$$

At first, this looks like good news for cointegration analysis: Assuming that measured wealth represents ownership of the capital stock k, there is a linear relationship between k and c and the coefficient on k should be the long-run marginal propensity to consume out of wealth.

But consider a related exercise. Suppose the world has two closed economies, A and B, that are identical to each other in every respect except time preference rates; the representative agent is more patient in one economy than in the other. In the more impatient economy, the equilibrium value of k will be lower, because there will be less capital relative to labor. If one were to postulate a coefficient measuring "the long-run MPC out of wealth μ" it would seem that such a thing could be measured from

$$c_a = 1 + \mu k_a \tag{3}$$
$$c_b = 1 + \mu k_b \tag{4}$$

so one could obtain an estimate from

$$\mu \;=\; \left(\frac{c_a - c_b}{k_a - k_b}\right). \tag{5}$$

The problem with this exercise is that there is no such object as μ. Instead, the relationship between consumption and wealth depends on the equilibrium interest rate, which differs across the two economies. In fact, assuming the same Cobb-Douglas aggregate production function $y = k^{\alpha}$ in both economies, it is easy to show that

$$rk \;=\; \left(\frac{\alpha}{1-\alpha}\right) \tag{6}$$

in both countries. But this implies that (5) will yield the answer $\mu = 0$, which is not a correct estimate of the MPC out of anything in either economy. In econometric terms, the problem is the endogeneity of the interest rate with respect to k.

What does this have to do with cointegration analysis? Well, if an individual economy spends a substantial amount of time in one equilibrium (say, the pre-1973 high-productivity-growth regime) and then a substantial amount of time in another equilibrium (say, a post-1973 slow-productivity-growth regime), one can make an identification of the pre-1973 economy with, say, economy A in the cross-section experiment, and the post-1973 economy with economy B, which gives the intuition that the cointegrating regression is no more valid than the cross-section regression.

The foregoing analysis concerned conintegrating estimates of the relation between the ratio of consumption to labor income and the ratio of wealth to labor income. A popular alternative is to seek a cointegrating relationship between the logs of consumption, labor income, and wealth. But a similar critique applies to that approach; indeed, because the relationship of the model in logs to the underlying consumption theory is looser (involving various approximations), my sense is that the log approach obscures the problem without in any way addressing it.

For the question at hand (the distinction between the effects of housing and stock wealth), the problems are even deeper, as a stable cointegrating vector would require not only a stable overall saving rate, but a lack of any trend in the proportion of wealth held in stock wealth and housing wealth. Again there is no reason from economic theory to expect this (especially in light of the massive changes in financial markets over the sample period), and from the historical experience of other countries as well as the U.S. there seems to be good reason to doubt it as an empirical proposition.

A seemingly-plausible response from advocates of cointegration analysis might be that the fact that the error from the cointegrating regression has predictive power for some of the variables contained in the cointegrating regression proves the validity of the cointegrating approach. However, a full-sample cointegrating regression gets to "pick" the saving rate (or cointegrating vector) that makes its fit to the model look best. This will tend to be a vector for which, *ex post*, the estimated errors tend to disappear over time. It is therefore neither surprising nor convincing to show that the error in the cointegrating regression has predictive power for either consumption or wealth changes. If there is even a rough sense in which the variables move together over time (e.g. if there is a cointegrating vector that gradually changes over time), the error must have power for one of them, and if there is more predictability to consumption growth (broadly speaking) than to wealth growth in high frequency data (because most high-frequency movements in wealth reflect stock market fluctuations), finding evidence for error correction in consumption is not persuasive that there is a stable cointegrating vector.

Another response might be that the foregoing arguments against cointegration rely too much on long-term general equilibrium relationships, and perhaps a cointegration approach makes more sense if we think of even a 40 year period like the one under examination in the U.S. as reflecting a partial equilibrium situation. But this only makes matters worse: There is a fundamental inconsistency between standard partial equilibrium consumption theory and the usual cointegrating regressions. As Hall 1978 taught us long ago, in the standard framework consumption should follow a random walk with a trend component that depends on tastes and interest rates but is not related to the underlying growth rate of labor income. Yet the organizing principle of the cointegrating approach is that consumption will reliably return eventually to a trend value, and so future changes in consumption are strongly predictable based on the cointegrating gap. While this idea is intuitively plausible (and, as I argue below, there is plenty of predictability to consumption growth), it is fundamentally inconsistent with the usual perfect foresight time-separable formulation of the consumption problem (at least in the short run).

It is interesting to speculate whether more complicated models with heterogeneous consumers subject to idiosyncratic shocks might modify this conclusion. However, as discussed below, there is a simpler framework that also leads to predictability in consumption growth over the medium term: a habit formation model. Using the basic equation for consumption dynamics in a habit formation model, I therefore developed an alternative approach to estimating the short- and long-run MPCs that does not depend on an assumption of the existence of a long-term stable cointegrating vector.

6.2 An Alternative

The starting point for the alternative approach is the Euler equation for consumption growth in a habit formation model of consumption.

Muellbauer 1988 proposes a utility function of the form

$$u(c,h) = \left(\frac{(c-\gamma_1 h)^{1-\rho}}{1-\rho}\right) \tag{7}$$

where γ_1 is a parameter capturing the importance of habits h to consumption (if $\gamma_1 = 0$, habits are irrelevant), and the habit stock evolves according to

$$h_{t+1} = (1-\lambda)h_t + \lambda(c_t - h_t) \tag{8}$$

and Dynan 2000 shows, using an approximation to the Euler equation, that the Euler equation to a shock to wealth for such a specification can be approximated by

$$\Delta \log C_{t+1} = \gamma_0 + \gamma_1 \Delta \log C_t + \epsilon_{t+1}, \tag{9}$$

so that the importance of habits can be estimated from the serial correlation parameter in consumption growth.

However, suppose our actual empirical measure of consumption contains either some transitory measurement error, or transitory elements of spending (e.g. weather-related fluctuations) that are not incorporated in the theory that leads to (9).

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

$$\Delta \log C^*_{t+1} = \beta_0 + \beta_1 \Delta \log C^*_t + \zeta_{t+1}. \tag{10}$$

Sommer 2001 points out that if the transitory element of consumption spending is confined to a quarter (which he argues is a reasonable assumption), while the process for true consumption follows (9) (possibly with $\gamma_1 = 0$, 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 (10) on NIPA data will yield an estimate of β_1 that is a downward-biased estimate of the habit parameter γ_1, where the size of the bias will depend on the magnitude of the transitory component of expenditures.

In the absence of further complications, γ_1 could be obtained by estimating an ARMA(1,1) model for consumption growth. Sommer does this, and obtains a highly statistically signficant estimate of the MA(1) coefficient, measuring about -0.3 to -0.5, with a sizable standard error, implying that the pure AR(1) estimation of (10) yields a strongly downward-biased estimate of the serial correlation of consumption growth. The ARMA(1,1) estimation yields an estimate of γ_1 in the neighborhood of 0.7-0.8. (These results are reproduced in the programs associated with this document).

However, if there is time aggregation as well as measurement error (as surely there is), matters are more complex, and the process for consumption can no longer be shown to follow a simple ARMA formula. Furthermore, ARMA coefficient estimates can be shown to be biased if the true process is more complex. However, Sommer shows that instrumental variables estimation using instruments dated $t-2$ or earlier should largely overcome these problems, and he presents IV estimates that again 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

$$\Delta \log C_t = Z_{t-1}\eta + \nu_t \tag{11}$$

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.

The results presented in tables 1 and 2 are based largely on the Sommer methodology. However, there are several features of the question that require some modification to his approach.

First, the ultimate goal here is to obtain an estimate of the marginal propensity to consume out of wealth. But (9) 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 were 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, because if the null hypothesis is equal MPC's 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 as each other in every period.

There is a simple solution to these problems, however, which is to use, rather than wealth growth, the ratio of changes in wealth to an initial level of consumption (because we will later be using variables with lags up to a year, the "initial" level here is defined as consumption 5 quarters before the current quarter). That is, if we define

$$dC_t = (C_t - C_{t-1})/C_{t-5} \tag{12}$$

$$\partial W_{t-1} = (W_{t-1} - W_{t-2})/C_{t-5} \tag{13}$$

and so on (where we are using the ∂ symbol for one-quarter changes in wealth because later we will be defining dW as a weighted average of a year's worth of wealth changes), then a first-stage regression of the form

$$dC_t = \alpha_0 + \alpha_1 \partial W_{t-1} \tag{14}$$

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

$$dC_t = \alpha_0 + \alpha_1 \partial W^S_{t-1} + \alpha_2 \partial W^N_{t-1} \tag{15}$$

yields directly comparable estimates of relative MPCs.

Regressions of the form (14) or (15) pass all the standard tests of instrument validity and therefore justify estimation of an IV equation of the form

$$dC_{t+1} = c_0 + \rho dC_t + \xi_{t+1}. \tag{16}$$

where c_0 is an unimportant constant.

The results of such IV regressions are what are reported in table 2. Except for the first row, all of the equations are overidentified, and all pass the standard overidentification tests. Furthermore, in experiments not reported here (but available from the author), 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 dC_t (e.g., an $\bar{R}^2$ of 0.1 or more) generates a highly statistically significant estimate of the ρ coefficient. Furthermore, the estimate of ρ tends to be larger the better is the performance of the first-stage regression. (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). The coefficient estimates in rows (1) and (3) of table 1 are precisely the coefficients of such an equation, and rows (2) and (4) report the results when the unemployment expectations and the Fed funds rate are added to the regression.

In order to focus on the bottom line for aggregate demand, the memo reports the results when the measure of C is total consumption expenditures. However, qualitatively identical results also hold for spending on nondurables alone, or nondurables and services, which are more theoretically appropriate measures of spending.

Given an initial (current-quarter) MPC out of wealth of μ and a serial correlation coefficient ρ for dC, 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

$$MPC^{LR} = \left(\frac{\mu}{1-\rho}\right). \tag{17}$$

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

My interpretation of the econometric object I call the "long-run MPC" in the memo 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 I am 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 my results reflect dynamics over a much shorter horizon.

Returning to the main thrust of how to estimate the "long run MPC," the simplest way to proceed 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 $\alpha = \rho\mu$ implies that the long-run MPC could have been estimated from

$$MPC^{LR} = \left(\frac{\alpha}{\rho(1-\rho)}\right). \quad (18)$$

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

Unfortunately, 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 or exclusion of other instruments (both the sentiment and Fed funds measures that I focus on in the memo, and alternative instruments like the lagged change in the unemployment rate) for me to be comfortable reporting them directly (thought not enormously sensitive - they typically generated long-run MPC's between 0.02 and 0.1). 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 ρ, assuming independent shocks to wealth from quarter to quarter we should have:

$$\Delta C_{t+1} \approx \mu\rho(\Delta W_t + \rho\Delta W_{t-1} + \rho^2\Delta W_{t-2} + \rho^3\Delta W_{t-3}) + \eta_{t+1}. \quad (19)$$

Now define

$$dW_t = (\Delta W_t + \rho\Delta W_{t-1} + \rho^2\Delta W_{t-2} + \rho^3\Delta W_{t-3})/C_{t-4} \quad (20)$$

and since similarly $dC_{t+1} = (C_{t+1} - C_t)/C_{t-4}$ this leads to an approximate equation for dC and dW of the form

$$dC_t = c_0 + \alpha_1 dW_{t-1}. \quad (21)$$

Under the assumption that the dynamic model of consumption is right, the coefficient estimate on dW_t should (among other things) be the immediate (first-quarter) MPC out of an innovation to wealth, as claimed in the main text.

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

$$MPC^{LR}_n = \left(\frac{\alpha_n}{\rho(1-\rho)}\right). \quad (22)$$

for the α corresponding to the respective measure of wealth.

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

1. Estimate (16) by IV, generating the estimate of ρ reported in table 2

2. Construct the estimate of dW as per (20)

3. Estimate (21) or the corresponding equation for the other instrument sets, yielding the estimate of the short-run MPC contained in 1

4. Construct the estimate of the long-run MPC for table 2 via (22)

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, but 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.

Figure 1: Real Housing Wealth Across Countries

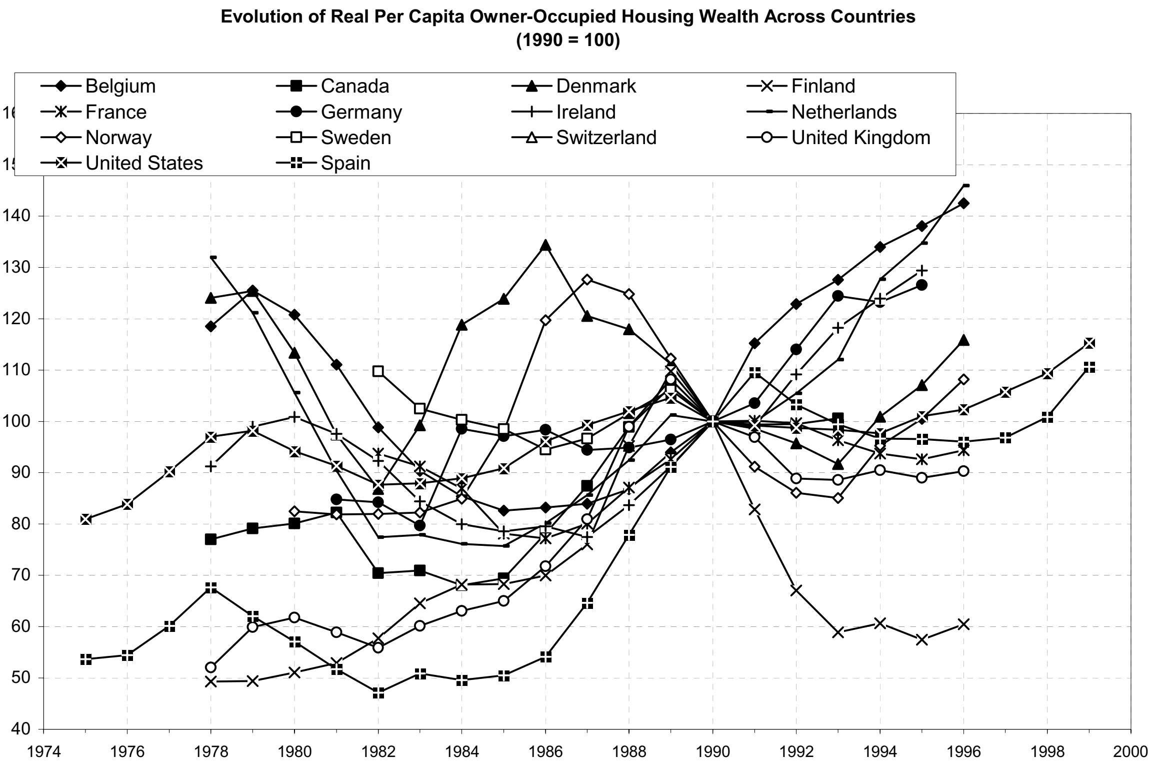

Source: Case, Quigley, and Shiller (2003), Figure 2.

References

M. J. Brennan and Y. Xia. Tay's as good as cay. *Manuscript, the Wharton School, University of Pennsylvania*, 2002.

John Y. Campbell and N. Gregory Mankiw. Consumption, income, and interest rates: Reinterpreting the time-series evidence. In Olivier J. Blanchard and Stanley Fischer, editors, *NBER Macroeconomics Annual, 1989*, pages 185–216. MIT Press, Cambridge, MA, 1989.

Glenn Canner, Karen Dynan, and Wayne Passmore. Mortgage refinancing in 2001 and early 2002. *Federal Reserve Bulletin*, 88:469–481, December 2002.

Karl E. Case, John M. Quigley, and Robert J. Shiller. Comparing wealth effects: The stock market versus the housing market. *Institute of Business and Economic Research Working Paper No. W01-004*, May 2003.

Morris Davis and Michael Palumbo. A primer on the economics and time series

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 $dW^S = dW^N$	$\bar{R}^2$
0.0173*** (0.0048)					(Assumed)	0.130
0.0076*** (0.0024)			0.086*** (0.034)	-0.655*** (0.200)	(Assumed)	0.280
	0.0163*** (0.0049)	0.0352*** (0.0116)			0.140 (Accepted)	0.135
	0.0075** (0.0028)	0.0149* (0.0077)	0.081*** (0.029)	-0.671*** (0.157)	0.351 (Accepted)	0.281

Notes: Sample period is 1960q3-2003q3 (the full sample for which all data were available). 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 0.0173 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; for details see the appendix. 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 8 lags.

Table 2: Consumption Growth Momentum and the Long Run MPC

$$dC_{t+1} = c_0 + \rho E_{t-1}[dC_t] + \zeta_{t+1}$$

Variables used to forecast $E_{t-1}[dC_t]$	Consumption Growth Momentum Coefficient ρ	Implied Long-Run MPC out of		
		Total W	Stock W^S	Nonstock W^N
W	0.57^{**} (0.23)	0.070		
W, MU, FF	0.83^{***} (0.11)	0.053		
W^S, W^N	0.49^{**} (0.19)		0.065	0.141
W^S, W^N MU, FF	0.78^{***} (0.10)		0.044	0.088

Notes: Sample period is 1960q3-2003q3 (the full sample for which all data were available). Standard errors are in parentheses. {*,**,***}=Statistical significance at {10,5,1} percent. The long-run MPC's are calculated from the formula $\alpha_n/\rho(1-\rho)$ 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. One subtlety: The measure of the change in wealth used for the regressions is the ∂W measure defined in the appendix, as this can be measured without an estimate of ρ, unlike the dW measures used in the previous table.

econometrics of wealth effects. *Federal Reserve Board Finance and Economics Discussion Papers 2001-09*, 2001.

Morris A. Davis and Jonathan Heathcote. The price and quantity of residential land in the united states. *Manuscript, Board of Governors of the Federal Reserve System*, 2004.

Richard Disney, Andrew Henley, and David Jevons. House price shocks, negative equity and household consumption in the uk in the 1990s. *Mimeo, University of Nottingham*, 2002.

Nikola Dvornak and Marion Kohler. Housing wealth, stock market wealth and consumption: A panel analysis for australia. *Reserve Bank of Australia Research Discussion Paper 23*, 2003.

Karen E. Dynan. Habit formation in consumer preferences: Evidence from panel data. *American Economic Review*, 90(3), 2000.

Gary V. Engelhardt. House prices and home owner saving behavior. *Regional Science and Urban Economics*, 26:313–336, 1996.

Jeffrey C. Fuhrer. An optimizing model for monetary policy: Can habit formation help? *American Economic Review*, 90(3), June 2000.

Xavier Gabaix and David Laibson. The 6d bias and the equity premium puzzle. *Manuscript, The Massachusetts Institute of Technology*, 2001.

Nathalie Girouard and Sveinbjrn Blndal. House prices and economic activity. *OECD Economic Department Working Paper 279*, 2001.

Richard K. Green. Stock prices and house prices in california: New evidence of a wealth effect. *Regional Science and Urban Economics*, 32:775–783, 2002.

Jaehoon Hahn and Hangyong Lee. On the estimation of the consumption-wealth ratio: Cointegrating parameter instability and its implications for stock return forecasting. *Manuscript, Columbia University*, 2001.

Robert E. Hall. Stochastic implications of the life-cycle/permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 96:971–87, 1978.

Hilary W. Hoynes and Daniel L. McFadden. The impact of demographics on housing and nonhousing wealth in the united states. In Michael D. Hurd and Yashiro Naohiro, editors, *The Economic Effects of Aging in the United States and Japan*. University of Chicago Press for NBER, Chicago:, 1997.

G. Donald Jud and Daniel T. Winkler. The dynamics of metropolitan housing prices. *Journal of Real Estate Research*, 23:29–45, 2002.

F. Thomas Juster, Joseph Lupton, James P. Smith, and Frank Stafford. The decline in household saving and wealth effect. *Mimeo, University of Michigan*, 2001.

Andreas Lehnert. Housing, consumption, and credit constraints. *Manuscript, Board of Governors of the Federal Reserve System*, December 2003.

Laurence Levin. Are assets fungible? testing the behavioral theory of life cycle savings. *Journal of Economic Behavior and Organization*, 36:59–83, 1998.

Alexander Ludwig and Torsten Sløk. The relationship between stock prices, house prices, and consumption in oecd countries. *Topics in Macroeconomics*, 4(1):Article 4, 2004.

John N. Muellbauer. Habits, rationality, and myopia in the life cycle consumption function. *Annales d'économie et de Statistique*, 9:47–70, 1988.

Misuzu Otsuka. Household portfolio choice with illiquid assets. *Manuscript, The Johns Hopkins University*, 2003.

Lawrence Slifman and Carol Corrado. Decomposition of productivity and unit costs. *American Economic Review*, 89(2):328–332, 1999.

Martin Sommer. Habits, sentiment, and predictable income in the dynamics of aggregate consumption. *Johns Hopkins University Department of Economics Working Paper Number 458*, 2001. http://www.econ.jhu.edu/pdf/papers/wp458.pdf.

Richard Thaler. Anomalies: Saving, fungability, and mental accounts. *Journal of Economic Perspectives*, 4:193–206, 1990.