

Does Stock Market Wealth Matter for Consumption?

Karen E. Dynan
Federal Reserve Board
Stop 93
Washington, DC 20551
(202) 452-2553
kdynan@frb.gov

Dean M. Maki
Putnam Investments
One Post Office Square
Boston, MA 02109
(617) 760-8616
Dean_Maki@putnaminv.com

May 2001

Abstract

This paper explores the household behavior that underlies the link between wealth and consumption at the aggregate level. One possibility is that changes in wealth directly *cause* changes in consumption through their effect on households' contemporaneous budget sets; another possibility is that they merely *predict* changes in consumption because they signal changes in future income. Previous attempts to assess the relative importance of these "direct" and "indirect" channels have yielded indeterminate results. Based on analysis of household-level data from the *Consumer Expenditure Survey*, we find that direct wealth effects begin to show up relatively quickly and continue to boost consumption growth for a number of quarters, in line with aggregate estimates. In contrast, we find that the indirect wealth channel is not an important determinant of consumption growth. We also estimate that an additional dollar of wealth leads households with moderate securities holdings to increase consumption between 5 cents and 15 cents, with the most likely gain in the lower part of this range.

Much of this paper was completed while Maki was an economist at the Federal Reserve Board. We wish to thank Byron Lutz, Shital Patel, and Richard Saouma for excellent research assistance and Chris Carroll, Morris Davis, Doug Elmendorf, Eric Engen, Michael Kiley, Spencer Krane, David Laibson, Michael Palumbo, Jonathan Parker, Martha Starr-McCluer, Jon Skinner, and participants at the NBER Summer Institute for helpful discussions. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Board or its staff.

1. Introduction

The dramatic climb in stock prices in the latter half of the 1990s has generated much discussion of whether "wealth effects" can explain the boom in consumer spending during these years. The simplest form of analysis is a graph of the personal saving rate and the ratio of household wealth to income such as that shown in Figure 1. These variables are strongly negatively correlated, and that pattern held up in the late 1990s as the wealth-income ratio soared to an unprecedented level and the saving rate plunged. More formally, econometric specifications of aggregate consumption such as that included in the Federal Reserve Board's FRB/US model (Brayton and Tinsley, 1996) generally show that an additional dollar of stock market wealth raises the level of consumer spending by 3 to 5 cents, with the effect emerging gradually over several years. Simple back-of-the-envelope calculations suggest that this sort of relationship can explain most of the decline in the saving rate since the mid-1990s.¹

Yet, relatively little is known about the household behavior underlying these findings. Budget constraints imply that an increase in a household's wealth must eventually raise its consumption. If this response emerges relatively quickly, it could explain the aggregate relationship through what we term the "direct" channel. However, a number of authors have noted that the response might occur with a substantial lag, owing to uncertainty about the permanence of a given stock market movement or a real or perceived (through a mental accounting framework) inaccessibility of stocks held indirectly through pensions and the like. Indeed, the lag could be so great that the wealth effect does not show up in the consumption of current equity owners but rather is passed on to future generations via larger bequests.² If lags

¹ If $\Delta C^* = -\Delta S^* = mpc \Delta W$, where C^* is target consumption, S^* is target saving, W is wealth, and mpc is the marginal propensity to consume out of wealth, then $\Delta(S/Y)^*$ is roughly equal to $mpc \Delta(W/Y)$, where Y is income. The wealth-income ratio rose by 1.4 over the six years ending in 2000:Q4. Assuming typical lags for consumption to adjust to the target and an mpc between 0.03 and 0.05, the wealth-driven decline in the saving rate over this period would be between 4 percentage points and 7 percentage points, as compared with an actual decline of 7-1/4 percentage points.

² The ultimate impact of a given increase in stock market wealth on a household's consumption depends on its source. An increase related to higher expected profits moves out the budget constraint, while an increase related to a reduction in the discount rate for profits changes the slope of the budget constraint. As with most aggregate estimates, our analysis simply

are long, the aggregate relationship between stock market wealth and consumption might arise because changes in stock prices predict changes in future income, or because of some other link; we refer to such effects as the “indirect” channel.

These alternative explanations for the aggregate relationship were first emphasized by Poterba and Samwick (1995) and are reviewed by Poterba (2000). The existing micro-data literature has provided fairly limited evidence on the relative roles of these channels. Some micro-data studies of the consumption CAPM (Mankiw and Zeldes, 1991, Attanasio, Banks and Tanner, 1998, Vissing-Jørgensen, 1999, and Brav, Constantinides and Geczy, 1999) find that the spending of stockholders is more highly correlated with stock market returns than that of nonstockholders, which supports a direct effect. However, Poterba and Samwick (1995) introduce other tests and conclude that the balance of evidence points to a small role for direct effects. Parker (1999a) and Juster, Lupton, Smith and Stafford (1999) both find that spending appears responsive to wealth at the household level, but neither paper can pinpoint the response as occurring in the time frame necessary to explain the macro relationship. Finally, Starr-McCluer's (2000) analysis of qualitative evidence from the University of Michigan's SRC Survey of Consumers suggests that the spending of stockholders is only modestly affected by changes in wealth, and Otoo (1999) finds that the correlation between stock prices and consumer sentiment does not vary by stock ownership. In sum, although there are strong theoretical arguments for direct wealth effects, the empirical evidence in hand does not decisively reject the view that "in the short run, at least, the stock market is a sideshow for consumption" (Shleifer, 1995, p. 359).³

estimates the average response of consumption to wealth over time. Households might also choose to reduce their labor supply and consume more leisure rather than consuming more goods. Some have pointed to wealth effects as an explanation of why the extraordinary availability of jobs in recent years has not led to a larger increase in the labor force participation rate, but there is limited empirical evidence on this point (see Cheng and French, 2000).

³ In addition to the micro-data studies discussed above, Maki and Palumbo (2001) combine aggregate Flow of Funds data with information from the *Survey of Consumer Finances* to create cohort estimates of saving by income and education groups and find that the decline in saving in the 1990s was concentrated among groups that experienced the largest increase in wealth. This result is consistent with the direct view of wealth effects.

The limited evidence on the household-level underpinnings of wealth effects partly reflects the lack of a good data set with which to explore the question. The *Consumer Expenditure Survey* (CE) has very detailed (albeit noisy) data on household expenditures, but much less data on assets and liabilities. The *Survey of Consumer Finances* (SCF) provides comprehensive data on household balance sheets at three-year intervals but contains essentially no information about consumption except what one can infer from the value of housing and vehicles. The *Panel Study on Income Dynamics* (PSID) collects information about participants' wealth every five years, but it also has little data on consumption aside from food expenditures.

This paper studies wealth effects using data from the CE. We believe that its shortcomings in terms of wealth data and other features are outweighed by several important advantages. First, comprehensive measures of consumption are necessary to determine the quantitative importance of wealth effects; extrapolating results for a few specific components of overall spending can be very misleading, as the elasticities of different goods could be quite different. Second, the quarterly frequency of the CE panel allows us to explore a time frame that is much more relevant to the timing of the macro relationship than surveys conducted less frequently like the PSID or SCF. Third, because CE interviews have been conducted continuously since the early 1980s (with new households replacing those that have completed their panel of five interviews), we can explore a richer set of stock market movements than with other surveys.

To preview our results, we find that direct wealth effects begin to show up relatively quickly and continue to boost consumption growth for a number of quarters, consistent with the lags in econometric models of the relationship between aggregate consumption and the stock market.⁴ In contrast, we find that indirect wealth effects are not an important determinant of consumption growth. For households with moderate levels of securities holdings, we estimate that every additional dollar of stock market wealth raises consumption between 5 cents and

⁴ Ludvigson and Steindel (1999) raise some questions about the short-term dynamics of the aggregate relationship, but Davis and Palumbo (2001) show that most reasonably specified aggregate consumption functions imply a gradual adjustment of household spending to capital gains. Kiley (2000) explores the simultaneity in the determination of aggregate consumption, income and wealth; he finds that the marginal propensity to consume out of “exogenous” wealth changes is a few cents per dollar.

15 cents, with the most likely gain in the lower part of this range. Because this response is larger than most aggregate estimates, it appears that households with high levels of securities holdings may have smaller responses to wealth gains. We discuss several explanations for such heterogeneity.

We provide a brief review of related literature in section 2 and discuss the data in section 3. Section 4 describes and provides results for our first empirical approach, which adapts a framework first used by Mankiw and Zeldes (1991) to explore the roles of direct and indirect wealth effects. Section 5 presents the second empirical approach, in which we calculate a household-specific measure of capital gains and use it to estimate the marginal propensity to consume. In section 6, we discuss conclusions and extensions.

2. Literature Review

Tests of the consumption CAPM using household-level data have shown that the consumption of stockholders is more highly correlated with excess returns on the stock market than the consumption of non-stockholders, which suggests at least some role for the direct channel. Mankiw and Zeldes (1991) discover this relationship, using annual observations on food consumption from the PSID. Attanasio, Banks, and Tanner (1998) confirmed the result for a broader measure of consumption, observed at a quarterly frequency, from the UK Family Expenditure Survey, and Vissing-Jørgensen (1999) and Brav, Constantinides, and Geczy (1999) produced similar findings using the CE. We believe that the Mankiw-Zeldes framework is a good starting place, and our first empirical approach modifies and extends that technique to more directly confront the question of interest in this paper. One innovation is that we add lags of stock returns to allow for gradual adjustment. In addition, our sample includes the period of rapid stock price growth in the latter part of the 1990s, which may allow for better estimates of the relationship between stock prices and consumption than produced by the other studies, whose samples ended earlier.

Two recent papers estimate the marginal propensity to consume out of wealth using panel data. Parker (1999a) constructs a proxy for nondurables and services consumption in the PSID by combining food consumption and demographic variables using weights estimated from CE data. He finds a significant positive relationship between consumption and wealth, but because

identification is based on cross-sectional variation in levels, Parker's findings have no implications for whether a direct wealth effect occurs in the short run. Juster, Lupton, Smith, and Stafford (1999) use data from the PSID wealth modules to examine the relationship between changes in so-called "active" saving and capital gains. They estimate a surprisingly large 17-cent marginal propensity to consume out of each dollar of additional stock market wealth. However, the results have limited implications for direct wealth effects at a one- to two-year frequency because the modules are five years apart. Moreover, because capital gains are constructed essentially as the difference between the change in total wealth and active saving, measurement error in active saving may be boosting their estimate of the marginal propensity to consume; for example, if active saving were overstated, capital gains would be understated, imparting a spurious negative correlation of the two variables.

Two papers based on the Michigan SRC Survey of Consumers also contain results related to the importance of the direct wealth channel. Starr-McCluer (2000) studies responses to qualitative questions in 1997 about the effect of wealth gains on spending. While her evidence supports a modest direct wealth effect, the fact that the vast majority of stockholders reported "no change in the amount of money [they] spend as the result of the trend in stock prices during the past few years" leaves open the possibility that the spending of many households is unaffected by capital gains. Otoo (1999) finds no significant difference between the correlations of stock market movements with the sentiment of stockholders and the sentiment of nonstockholders. Her results are thus consistent with the view that households use stock prices as a leading indicator of future economic activity, although they do not rule out a direct wealth effect.

Finally, Poterba and Samwick (1995) test for direct wealth effects by examining whether stock prices influence the share of aggregate spending devoted to goods that are consumed disproportionately by high-income households. Of the "luxury" goods they identify — new cars, education, hotel and motel spending, domestic services, and entertainment — they find that only the share of new cars is significantly positively correlated with stock price movements, and argue further that even this relationship probably arises because the stock market forecasts consumer demand.

3. The Consumer Expenditure Survey

The CE is a quarterly survey of about 5000 households that has been conducted since 1980. Each household is interviewed five times at three-month intervals before being rotated out and replaced with a new household. The public-use data sets include information from the second through fifth interviews, in which households are asked for very detailed data about their expenditures over the preceding three months. In the fifth interview, households are asked a few questions about assets and liabilities, relating to both current levels and changes over the preceding year. Information about the income and demographic profiles of respondents is also gathered, primarily during the second and fifth interviews.

Following Parker (1999b), our primary measure of consumption includes all nondurables and services categories of aggregate consumption from the National Income and Product Accounts (NIPA) except for housing, health care, education, and charitable contributions. The housing and medical expenditures included in the CE — in particular, mortgage payments and out-of-pocket health costs — do not correspond to the imputed service flow from housing (constructed from rental values) or total health spending included in NIPA consumption, so they are not useful for linking household behavior to aggregate estimates. We exclude education because it may be viewed by households more as investment than consumption. Whether charitable contributions represent consumption is more debatable, but we found that including this component made little difference to our results. For each household, we construct one reading of consumption from each of their interviews. We follow earlier authors (e.g. Parker, 1999b; Souleles, 1999a) in controlling for normal seasonal variation in spending by including a set of dummies that indicate the month corresponding to the final interview.

As noted above, balance sheet information in the CE is limited. In the fifth interview, households are asked: “Did you (or any members of your [consumer unit]) own any securities, such as stocks, mutual funds, private bonds, government bonds, or Treasury notes on the last day of last month?” Respondents who answer “yes” are asked for the estimated value of all such securities on the last day of the previous month and for how this amount compares with the value of such securities one year earlier. Respondents are then asked for the total purchase price (including broker fees) of any securities purchased during the past twelve months and for the amount received from sales (after subtracting broker fees) of any securities over the past twelve

months.

Because there are no questions specifically about equity holdings, our analysis requires assumptions about the fraction of “securities” that reflects stock holdings and whether households include the value of defined contribution pension plans in their responses. In addition to these obstacles, two specific features of the data set limit information about the richest households. First, Sabelhaus (1998) showed that the CE sample appears to do a poor job of covering households at the top of the income distribution, where much of equity wealth is held. Second, wealth and other variables are topcoded, meaning that if the variable exceeds a certain critical value, it is flagged and set equal to some pre-determined value. This practice is most problematic prior to 1995, when both the critical and topcoded values for securities holdings equaled \$100,000. Beginning in 1996, the critical value for securities holdings was raised substantially (to \$500,000 for that wave, followed by further increases in subsequent waves), and the topcoded value was set equal to the mean across all topcoded observations. This change reduced the fraction of households with nonzero securities holdings who had topcoded values from about 16 percent right before the change to an average of 2 percent since the change. While this implies that the post-1995 data are far superior in this respect, excluding the earlier years yields a sample that is generally too small to obtain precise results.

Despite these problems, a number of authors have concluded that the CE balance sheet data contain useful information. Sabelhaus (1993) shows saving rates derived from the CE change in net worth “are more reasonable than the residual measures commonly used” (p. 332). Attanasio (1994) concludes that “the main features of the financial asset data [in the CE] are similar to those found in other surveys” (p. 63). Souleles (1999b) considers the correlation between securities purchases and various measures of hedging motives and finds a pattern that is consistent with theoretical predictions. Maki (2001) finds that the CE credit card and auto loan data match up well with data from the *Survey of Consumer Finances*.

Our sample excludes households with missing or invalid readings for the wealth variables. Many of the public-use CE data sets from the early and mid-1990s contain miscoded data, listing households with invalid or missing responses to the securities as having valid

responses of \$1.⁵ Since households are extremely unlikely to have actual holdings of \$1, we also drop the observations equal to this amount. Topcoded households can be included in our first empirical exercise, where wealth is used only to classify households as stockholders or nonstockholders, but not in our second exercise, where we calculate and analyze a household-level measure of capital gains. For consistency, we focus on the same sample of non-topcoded households in both exercises, but note any place in the first exercise where including the topcoded households makes a difference. To keep the sample uniform in the face of the 1996 change in topcoding practices, we drop households in later waves whose wealth would have been topcoded under the pre-1996 criterion.⁶

Our sample is drawn from CE data files corresponding to the period 1983:Q1 (earlier data had bigger problems with data quality) through 1999:Q1 (the latest date for which CE data are available). In addition to the restrictions noted above, we exclude households with incomplete income responses, households that changed size or marital status, households lacking either a second or fifth interview, households living in student housing, households with more than one consumer unit, and households with implausibly small food consumption (less than \$100 for any three-month period). Finally, following Zeldes (1989a) and much of the rest of the household consumption literature, we drop households that had extremely large (absolute value) changes in consumption between the second and fifth interview (above the 99th percentile or below the 1st percentile). We deflate all nominal variables by a chain price index corresponding to the consumption measure, constructed by dividing a similar nominal NIPA aggregate by its real (chain-weighted) counterpart.

⁵ The economists of the Bureau of Labor Statistics who construct the CE confirmed that the error occurred in all quarterly data sets from 1991:Q2 through 1995:Q4 and from the first half of 1997.

⁶ We recognize that this criterion is somewhat arbitrary because the critical value is fixed in nominal terms. However, less arbitrary topcoding criteria — such as dropping the same fraction of households from each wave or imposing an inflation-adjusted critical value — result in a significant loss of the relevant variation. We experimented with these alternative criteria and generally found similar point estimates but higher standard errors.

4. Direct Wealth Effects and the Consumption of Stockholders and Nonstockholders

Our first approach adapts the Mankiw-Zeldes framework to our question of interest. The basic idea is to divide the sample by stock ownership and examine the correlation between households' consumption growth and aggregate stock market returns; if stock market wealth has no direct effect on consumption, this correlation should be no higher for the stockholders. One advantage to this approach is that we rely on the noisy CE securities variable only to classify households as stockholders or nonstockholders.

Setup

We implement the framework by estimating the following equation for stockholders and nonstockholders separately:

$$\Delta \ln C_{it} = \alpha_0 + \sum_{j=0}^N \beta_j \Delta \ln Wil_{t-j} + \gamma X_{it} + \epsilon_{it}, \quad (1)$$

where $\Delta \ln C_{it}$ is the change in the log of real consumption, $\Delta \ln Wil_t$ is the change in the log of the real value of the Wilshire 5000 index of stock prices, and X_{it} is a vector of control variables.⁷ Letting t pertain to the month of the fifth (final) interview, C_{it} refers to real consumption in months t , $t - 1$, and $t - 2$ (i.e. the reference period for the final interview), and $C_{i,t-3}$, $C_{i,t-6}$, and $C_{i,t-9}$ refer to consumption covered by the fourth, third, and second interviews, respectively. The change in consumption, $\Delta \ln C_{it}$, is the log-difference between consumption in the fifth and second interviews, and the change in the Wilshire, $\Delta \ln Wil_t$, is the log-difference between the average values in the three months covered by the fifth interview and the three months covered by the second interview. This timing scheme is illustrated in Figure 2.

Poterba and Samwick (1995) tested for direct wealth effects by looking at the contemporaneous correlation between aggregate stock returns and food consumption growth in the PSID and found a stronger correlation for stockholders than nonstockholders. Similar results were obtained using a proxy for nondurables and services consumption created by combining selected PSID variables related to consumption with weights estimated from the CE.

⁷ We use the Wilshire 5000 because it was designed to capture the performance of the entire stock market.

Regressions used to estimate weights for such proxies typically explain a good deal of the cross-sectional variation in the level of consumption, but this does not imply that changes in the proxy will capture changes in consumption well. Of particular note, some of the PSID variables are quite durable (e.g. housing and cars) and therefore probably have very different time-series properties than nondurables and services consumption.

Vissing-Jørgensen (1999) and Brav, Constantinides, and Geczy (1999) have already used the CE data to demonstrate a stronger contemporaneous correlation between stock returns and CE consumption growth for stockholders than for nonstockholders. The contemporaneous correlation is appropriate for their topic — the consumption CAPM — but less useful for our purposes because the macro wealth effect appears to occur gradually over a period of several years. Therefore, we include lagged values of the change in stock prices in our specification. To be specific, since the “contemporaneous” value of the change in the Wilshire corresponds to a three-quarter change, the first lag corresponds to the change over the three quarters preceding the second interview, the second lag corresponds to the change over the three quarters preceding that, and so on — as shown in Figure 2.

Our set of control variables includes monthly dummies to allow for normal seasonal patterns in consumption (as noted above), and year dummies to allow for aggregate shocks unrelated to the stock market. We control for age, age-squared, and family size as tasteshifters, and we also control for a set of socioeconomic variables — income as well as dummy variables corresponding to educational attainment and whether the head of household is white — that have been found to be significantly correlated with consumption growth, possibly because they are correlated with underlying preferences (see Lawrance, 1991). Finally, our specifications all control for growth in earnings, which may affect consumption through its signal about permanent income or directly because of liquidity constraints or myopic behavior.

Results

Tables 1 through 3 present results from estimating equation (1). Following other authors, we split the sample in several ways: we first classify households as stockholders if they have real securities holdings greater than zero, then if they have real securities holdings greater than

\$1000, and finally if they have real securities holdings greater than \$10,000.⁸ In all cases, we classify households as nonstockholders if they have real securities holdings below the dividing line.

Table 1 includes just the contemporaneous change in stock prices and confirms previous findings using this framework. The estimated relationship between stock returns and the consumption growth of stockholders is positive, stronger when the group is limited to those with greater holdings, and statistically significant at or close to the 5 percent level. In contrast, stock returns appear to have a negative and statistically significant effect on the consumption growth of nonstockholders. This latter result seems anomalous — albeit consistent with Poterba and Samwick (1995) and Vissing-Jørgensen (1999) — but it fades substantially when we add lagged stock returns (see below).

Looking at the other estimates in table 1, the generally negative coefficients on age and sometimes positive coefficients on age-squared are consistent with the hump-shaped pattern of consumption over the life-cycle documented by Carroll and Summers (1991). However, the age variables and other socioeconomic variables are, for the most part, insignificant. Earnings growth has the expected positive sign for the full sample, but much less effect for stockholders than for nonstockholders. It may be that liquidity constraints make the consumption of lower-wealth households more responsive to income. Together, the explanatory variables account for only about 1 percent of the variation in consumption growth, consistent with other studies of consumption growth using household-level data.

Table 2 adds lagged changes in stock prices. For stockholders, the estimated response of consumption growth to contemporaneous stock returns is much larger when lags are included, and the lags themselves have positive estimated effects on consumption growth. Because stock returns are negatively serially correlated in this sample, specifications that omit the lags — as in table 1 — lead to downward-biased coefficients for the contemporaneous variable. Intuitively, households react in a given period to current *and* past wealth movements, and since the change

⁸ Since stocks most likely represent the majority but not all “securities” reported by CE households (as discussed in more detail below), we are effectively splitting the sample according to somewhat lower values of real stockholdings. The results do not change materially when we apply the splitting criteria to estimates of real stockholdings.

in the previous period generally had an opposite sign in this sample, the lagged response to past changes offsets some of the contemporaneous response to current changes. The coefficients on the first lags (corresponding to the change over the three quarters ending in the households' second interview) are about as large as those on the contemporaneous change, and significant at close to the 5 percent level for the most restrictive definition of stockholders and between the 5 percent and 10 percent level for the other definitions.⁹ The coefficient on the second lag is smaller and significant at the 15 percent level only for the most restrictive definition of stockholders; the coefficient on the third lag is also positive, but it is much smaller and statistically insignificant. The implied timing of the wealth effect is therefore much the same as that found in most macro models, with much of the impact occurring in the first couple of years and the effect trailing off after that.

The significance of the lagged terms for stockholders seems at odds with standard life-cycle/permanent income models, which generally imply that lagged information should not help to predict current changes in consumption (Hall, 1978). Macro models often justify using lagged information by assuming that behavior is characterized by habit formation (for example, see Fuhrer, 2000). Alternatively, households may adjust consumption relatively infrequently because of the costs associated with gathering information about changes in their portfolios and recalculating optimal consumption. Gabaix and Laibson (2000) show that infrequent adjustment at the household level will, under plausible assumptions, lead aggregate consumption to respond slowly to shocks to permanent income and the average across households (as we estimate here) to exhibit what looks like a gradual adjustment process.¹⁰

⁹ When topcoded households are included in the sample of stockholders, the coefficients on the contemporaneous Wilshire term are 15 to 30 percent smaller but still statistically significant (when only the contemporaneous term is included — as in table 1 — they are not statistically significant). The coefficients on the first lag are a bit larger, and all statistically significant at the 5 percent level or better.

¹⁰ Direct household-level evidence that could distinguish between the habit formation and delayed adjustment models is limited. One implication of habit formation is that one should observe positive serial correlation in consumption growth at the household-level; Dynan (2000) finds no such correlation for annual food consumption growth, suggesting that habit formation is unimportant. However, the lack of serial correlation in household-level consumption growth rates would be consistent with delayed adjustment since that model implies that households

For nonstockholders, adding lagged wealth changes attenuates the estimated negative contemporaneous relationship. While the point estimates remain below zero, they are smaller in absolute value and statistically insignificant. The point estimates for the lagged terms are positive — as one might expect if lagged stock price increases boost the consumption of nonstockholders via a multiplier effect — but they are all statistically insignificant.

The results in Table 2 could reflect indirect wealth effects rather than direct effects if the stock market were a better leading indicator of the wages of stockholders than nonstockholders. Table 3 explores this possibility. The CE does not contain measures of the hourly wages of respondents, but it does have information about total labor earnings. In the top panel, we replace consumption growth as the dependent variable with growth in household earnings. (This variable is dropped from the set of controls, but the other variables in the regressions are the same as in the previous tables.) The results provide no evidence that stockholders' labor earnings follow stock market movements more closely than earnings of nonstockholders. In fact, for the most restrictive definition of stockholders, the correlation between earnings and the stock market appears to be negative. Of course, earnings reflect not only the wage, but also the amount of labor supplied; as a result, a negative relationship between earnings and the stock market could be explained by households deciding to consume more leisure in response to wealth gains. In the lower panel of the table, the dependent variable is growth in hours worked by the head and spouse (where hours are defined as the product of weeks worked in the past year and hours "typically worked" in a week). There is no evidence that households reduce labor supply in response to current or recent capital gains, although the coefficient on $\Delta Wilshire_{t-27}$ for the most restrictive definition of stockholders suggests that a response might emerge after a number of quarters. Because the negative relationship between earnings and stock prices shown in the top part of the table has a different timing, it does not appear to be driven by changes in labor supply.

adjust infrequently to new information.

5. Estimating the Marginal Propensity to Consume out of Wealth

Turning now to the question of the magnitude of the marginal propensity to consume, we begin by considering the relationship between consumption and wealth implied by theory, then move on to our empirical implementation.

Theoretical Considerations

Assume first that the utility of household i in period t takes the CRRA form:

$$u(C_{it}) = \frac{1}{1 - \gamma} C_{it}^{1-\gamma}, \quad (2)$$

and that there is no uncertainty. Letting δ equal the rate of time preference, r the (constant) real rate of interest, T the household's final period, and W_{it} the sum of financial wealth and the present discounted value of future earnings, one can easily show that consumption in period t is given by:

$$C_{it} = \left(\frac{1 - \xi}{1 - \xi^{T-t+1}} \right) W_{it}, \quad (3)$$

where

$$\xi = \frac{1}{1 + r} \left(\frac{1 + r}{1 + \delta} \right)^{\frac{1}{\gamma}}. \quad (4)$$

The term multiplying W_{it} in equation (3) is the marginal propensity to consume out of wealth: the amount C_{it} would rise if wealth were to rise by \$1. If T is large relative to t , the implied marginal propensity to consume typically is small. For example, if $r = \delta = 0.03$ and $T - t = 60$, then the implied marginal propensity to consume is 0.035. Intuitively, the number is small because the household spreads the gain over the remainder of its life — it therefore raises consumption by just the annuitized value of the increase in wealth.¹¹

¹¹ Poterba (2000) finds that the implied marginal propensity to consume ranges between 2-1/2 percent and 10 percent for different values of r and $T - t$. Starr-McCluer (2000) calculates that the marginal propensity to consume for an average stockholder would be 0.033.

This model suffers from several shortcomings. To begin, it lacks features like habit formation or delayed adjustment which, as discussed previously, would produce a gradual response of consumption to wealth changes. Moreover, the model ignores uncertainty in order to derive a closed-form solution to the problem. Numerical simulations by Zeldes (1989b) and analytic results by Carroll and Kimball (1996) show that the introduction of uncertainty produces a concave consumption function: the marginal propensity to consume out of wealth is higher for households with lower resources. Parker (1999c) uses nonparametric methods to estimate the relationship between consumption and liquid assets for CE households, and finds empirical support for such concavity.

Thus, theory has several implications for our empirical analysis. First, the appropriate specification is in levels (rather than logs): a change in wealth in equation (3) produces a proportional change in the level of consumption. Second, in the absence of uncertainty, the marginal propensity to consume out of wealth would likely be several cents on the dollar. Third, in the presence of habit formation or delayed adjustment, lagged changes in wealth will affect consumption. Fourth, the marginal propensity to consume for lower wealth households would likely be higher in the presence of uncertainty than in the no-uncertainty case. Therefore, because our CE sample excludes the richest households, we might find a marginal propensity to consume that is larger than the average for all households and larger than that predicted by a model without uncertainty.

Identification Strategy

One approach is to use cross-sectional data and estimate a relationship like equation (3) in levels. Unfortunately, such cross-sectional identification would, at best, yield information only about the long-run marginal propensity to consume, as in Parker (1999a). In addition, unobservable variables such as differences in risk aversion or discount rates might vary systematically across the wealth distribution and contaminate estimation of the true causal relationship between wealth and spending.

Another possible approach is to construct synthetic cohorts from cross-sectional data; that is, to divide households in each cross-section between those that are likely to be stockholders and those that are not. Then we could compare changes in wealth and consumption

for stockholding households with changes for other households over time. The resulting estimates would not rely on cross-sectional variation, and could include longer lags of wealth changes than the relatively short panel in the CE would allow directly. Unfortunately, two features of the CE data militate against this approach.

The first problem is that so many households have topcoded securities holdings: by 1995, about 16 percent of all households with positive holdings are topcoded. As a result, average securities holdings for the total population and for wealthier groups will have serious inaccuracies, and changes in average holdings over time will also be distorted as more holders became topcoded. Attanasio (1994) demonstrates that conditional means can be constructed with topcoded data by estimating a cross-sectional density function and assuming that this density fits the unobserved tail of the distribution, but he notes that “the reliability of such an estimate hinges in a critical way on the parametric specification used” (p. 67). Reliability would be especially problematic when such a large percentage of households are in the distribution’s tail, and we do not attempt such an estimation strategy in this paper.

A second problem with the synthetic cohort approach is that the CE data do not show the dramatic fall in saving during the mid to late 1990s that is evident in the National Income and Product Accounts. Table 4 shows aggregate CE income, expenditures, and saving, both in dollars and as a percentage of the corresponding NIPA figures. Conceptual differences account for some of the average gap between the NIPA and CE figures, but are unlikely to account for changes in the gap over time. CE after-tax income has been a roughly constant fraction of NIPA after-tax income, but CE spending has fallen noticeably as a fraction of NIPA spending. As a result, the CE saving rate (where saving is defined as the difference between after-tax income and expenditures) rose over the 1990s, while the NIPA saving rate fell. Sabelhaus (1998) and Parker and McCully (2000) show that the decline in CE expenditures relative to the NIPA has been occurring since the 1980s. One possible explanation is that wealthier households are less conscientious about completing the lengthy CE questionnaire as their opportunity cost grows ever higher. In any event, any change in CE coverage over time that does not affect all types of households equally would distort a synthetic cohort analysis, and we do not pursue that technique.

Using the CE Panel

We believe that making use of the short panel in the CE is the best approach to estimating the marginal propensity to consume out of wealth. Because the CE is a true panel, we do not have to rely on cross-sectional variation in consumption levels and wealth for identification, and we can focus on households that do not have topcoded wealth.¹² Also, the downward drift in CE spending relative to the NIPA is unlikely to distort a given household's expenditures materially over the short panel.

A natural strategy is to regress the change in consumption between the second and fifth interviews on the change in the values of securities holdings. However, households that decide to increase consumption may liquidate part of their wealth to do so, and a simplistic regression would find a spurious negative correlation between wealth and consumption changes. Therefore, we estimate the relationship between the change in consumption and the *passive component* of changes in stock market wealth — capital gains, $capgn_{it}$:

$$\Delta C_{it} = \beta_0 + \sum_{j=0}^N \beta_j capgn_{i,t-j} + \gamma X_{it} + \epsilon_{it} . \quad (5)$$

We include various lags of capital gains, and X_{it} represents a vector of control variables, as before.

A straightforward way to construct $capgn_{it}$ is to take the reported change in securities holdings, subtract reported purchases, and add reported sales. Unfortunately, the reported change in securities may have more noise than signal, as a substantial fraction of households holding securities at the fifth interview reported no change in the preceding year — even in periods when aggregate stock prices climbed rapidly. For example, in the post-1995 sample, 30 percent of households with positive securities holdings reported no change in their holdings, as did 27 percent of households with holdings greater than \$5000.¹³ We suspect that the problem arises partly from households misinterpreting the question as asking about active changes to their

¹² We will therefore be estimating the average marginal propensity to consume for the group of non-topcoded households rather than that for the population as a whole.

¹³ Attanasio (1994) notes similar problems in the CE. Kennickell and Starr-McCluer (1997) study retrospectively reported portfolio changes using the 1983-1989 *Survey of Consumer Finances* panel and also find an overreporting of responses like “no change.”

portfolio: 7 percent of those reporting purchases or sales responded that the change in their holdings *exactly* equaled the difference between these two variables. Because of these problems, we do not use the reported change in securities in our analysis.

We instead form a measure of capital gains by combining households' final *level* of securities with data on its sales and purchases of securities, and with aggregate stock returns. The first step is to convert the value of "securities" to a value for *stocks*. Unfortunately, the CE question is somewhat vague about what sorts of assets should be included in securities holdings, asking for items "such as stocks, mutual funds, private bonds, government bonds or Treasury notes." To gain further insight into how households interpret the question, table 5 compares the fraction of CE households holding "securities" with fractions of households from the *Survey of Consumer Finances* (SCF) holding different groupings of assets that plausibly might be included in the CE responses. The CE shares are quite similar to the shares of SCF households holding stocks, bonds, and mutual funds, but much smaller than the SCF shares including defined contribution retirement assets, suggesting that CE households did not include these assets when responding to the securities question.¹⁴ In light of these results, we assume that stocks held by each CE household are:

$$stk_{it} = sec_{it} * eqshare_t, \quad (6)$$

where sec_{it} is reported securities holdings and where $eqshare_t$ equals aggregate stock holdings divided by the aggregate value of stocks, bonds, and mutual funds.^{15, 16}

We also use $eqshare_t$ to calculate net purchases of stocks for each household:

¹⁴ We do not include U.S. savings bond holdings in the SCF groupings (or in other calculations) because the CE asks about them separately.

¹⁵ The sensitivity of the results to this assumption will be explored below. An issue that is potentially more important than the exact definition of securities is that the return on the non-stock portion of securities may be positively correlated with the return on stocks. In this case, our method would understate the wealth gain associated with a given movement in the stock market, and our estimate of the marginal propensity to consume will be biased upward. We return to this issue below.

¹⁶ In calculating this share, we include both stocks held directly and stocks held indirectly in the numerator. Because the Flow of Funds data we use for the calculation are available only on a quarterly basis (or end-of-year for some series), we interpolated to obtain monthly values.

$$netpur_{it} = ((1 - 0.015)psec_{it} - (1 - 0.015)ssec_{it}) eqshare_t \quad (7)$$

where $psec_{it}$ corresponds to household i 's reported purchases of securities over the twelve months preceding the fifth interview, and $ssec_{it}$ corresponds to reported sales of securities over this period. We multiply purchases and sales by $(1 - 0.015)$ because households are asked to report purchases inclusive of brokerage fees and sales net of fees, and we assume that brokerage fees equal 1-1/2 percent of the value of the securities transacted.¹⁷ Assuming that all purchases and sales were at the midpoint of the twelve-month reference period for these questions, and that they earn the return on the Wilshire index on average, household i 's stock ownership in the absence of net purchases would have been:

$$stk'_{it} = stk_{it} - netpur_{it} \frac{Wil_t}{Wil_{t-6}}, \quad (8)$$

where Wil_t equals the average value of the Wilshire 5000 in month t .¹⁸

With these variables in hand, we can now create a measure of capital gains. To match the timing of the consumption change, we construct the capital gains between the three-month period corresponding to a household's second interview and the three-month period corresponding to a household's fifth interview. Denoting backward-looking three-month moving averages with an overbar, the average value of stocks during the consumption reference period for the second interview is:

¹⁷ Our results are not sensitive to reasonable variation in this assumption. Note that $eqshare_t$ is not ideal for parsing out the portion of net purchases related to stocks as it reflects the composition of the *level* rather than the *flow* of securities, but we cannot readily construct the latter concept.

¹⁸ To keep things simple, this formulation assumes that the reported value of securities for the fifth interview equals the average value for month t . Although respondents are literally asked about the value as of the last day of the preceding month, it seems reasonable to think that they are more likely to know the approximate range over which it has been fluctuating in recent times. Of course, the choice of a one-month average is somewhat arbitrary; however, changing our assumption to, for example, a two-month or three-month average makes little difference for our results.

$$\overline{stk}_{i,t-9} = \frac{1}{3} (stk_{i,t-11} + stk_{i,t-10} + stk_{i,t-9}) . \quad (9)$$

The CE does not report securities holdings in each of those months, but we can back out the information by assuming that households' portfolios earn the return on the Wilshire index on average, and by using our constructed measure of stockholding if there had been no net purchase. Specifically, we set:

$$\overline{stk}_{i,t-9} = stk'_{it} \frac{\overline{Wil}_{t-9}}{Wil_t} . \quad (10)$$

Finally, we calculate capital gains as:

$$capgn_{it} = \overline{stk}_{i,t-9} \left(\frac{\overline{Wil}_t}{\overline{Wil}_{t-9}} - 1 \right) + netpur_{it} \left(\frac{\overline{Wil}_t}{Wil_{t-6}} - 1 \right) . \quad (11)$$

Thus, we multiply average stockholdings for the three months corresponding to the second interview by the *net* return on the Wilshire between that period and the three months corresponding to the fifth interview; we then add the likely capital gains on net purchases, assuming that all sales and purchases were done at the midpoint between interview reference periods.

Lagged value of capital gains have timing similar to those in earlier regressions: the first lag corresponds to the passive change in stock holdings in the second interview from three quarters earlier, the second lag corresponds to the change over the preceding three quarters, and so on, as shown in Figure 3. We calculate the first lag of capital gains by constructing a value for stocks held in an earlier period and then multiplying by the net return on the Wilshire over the relevant period:

$$capgn_{i,t-9} = \overline{stk}_{i,t-18} \left(\frac{\overline{Wil}_{t-9}}{\overline{Wil}_{t-18}} - 1 \right) . \quad (12)$$

We calculate the second and third lags of the capital gains variable similarly. Unfortunately, we cannot remove net purchases between $t - 18$ and $t - 9$ (or earlier periods for the longer lags) because we do not have data on sales and purchases for those periods; thus, we are implicitly

assuming that net purchases equal zero for each household. This unavoidable assumption undoubtedly adds noise to our analysis: the capital gains of households who sold stocks (on net) over these periods will be understated, and the capital gains of households who purchased stocks will be overstated. As noted above, if we were to ignore net purchases during the period in which we are measuring the change in consumption, the mismeasurement of capital gains might bias our results. However, ignoring net purchases over periods preceding the consumption reference period should not be problematic as long as households that liquidate stocks to fund consumption changes do so fairly rapidly.

Results

The first column of table 6 shows estimates of equation (5) using the same set of control variables as in the earlier empirical exercise; the coefficients on these variables are not shown but are qualitatively similar to those in the first column of table 1. The coefficient on the contemporaneous capital gains term indicates that a capital gain of \$1 over the preceding nine months raises the consumption change over the same period by 9 cents, significantly different from zero at less than the 1 percent level. The coefficient on the first lag indicates that the subsequent nine-month change in consumption is higher by 5 cents, statistically significant at about the 6 percent level. The point estimate on the second lag of the capital gains variable suggests a further small rise in the next period, but neither this coefficient nor the small negative coefficient on the third lag can be statistically distinguished from zero. In the fifth row of the table, we show the sum of the coefficients on the contemporaneous capital gains variable and its three lags. This sum is the marginal propensity to consume — the total change in the level of consumption associated with a \$1 capital gain after roughly two years.¹⁹ The point estimate on the sum in the first column suggests a marginal propensity to consume of 0.132, statistically different from zero at less than the 1 percent level.

The other columns of table 6 show that these results are robust to changes in the control

¹⁹ Adding only the statistically significant coefficients (on the contemporaneous term and the first lag) has little effect on the implied marginal propensity to consume since the coefficients on the second and third lags roughly cancel out in all of the specifications. Dropping these lags from the regressions produces similar results.

variables. Excluding the first difference of earnings, excluding that variable and the socioeconomic variables, or adding dummy variables that interact year with education, all have little effect on the coefficients of interest. The last alternative is a crude way of controlling for changes in the returns to education over our sample period. If stock prices have been affected by the same forces that have changed the college wage premium over the past two decades, the estimated coefficients on capital gains might be inflated because the consumption change of stockholders (who tend to be more educated than the rest of the population) would be boosted by relative increases in their lifetime labor income. However, these data do not suggest any bias of this sort.

The estimates of the marginal propensity to consume in table 6 are several times larger than the figures typically estimated from aggregate data or implied by simple models of consumption. The disparity is even more striking when one considers that our consumption measure excludes durable goods, where some additional response to wealth gains would be expected. One possibility is that the non-topcoded households we study are not representative of the population as a whole; we discuss this issue later in the paper. Alternatively, our method for calculating the capital gains may have systematically understated the changes in wealth, which would overstate the marginal propensity to consume. Table 7 and table 8 explore the possibility of such a bias.

One issue is that our aggregate equity share reflects the holdings of nonprofit organizations (which are included as part of the household sector in the Flow of Funds accounts) and very wealthy households that are excluded from our sample. The top panel of figure 4 shows that the Flow of Funds equity share is generally lower than average equity shares in SCF samples that are restricted like our CE sample.²⁰ For the second column of table 7, we calculate

²⁰ Figure 4 excludes SCF households with holdings of stocks, bonds, and mutual funds greater than \$100,000. We cannot do this calculation for the 1983 SCF wave because that survey had insufficient detail to parse out the fraction of mutual funds represented by equities. Note that the SCF figure moves down between 1995 and 1998 in contrast with the upward movement of the aggregate series; this difference reflects both the restrictions we have placed on the SCF sample and the fact that the SCF number is the average of a ratio instead of the ratio of two aggregates. The ratio of the two aggregates for the unrestricted SCF samples shows the same upward trend seen in the Flow of Funds series.

capital gains based on SCF equity shares, using the value from the 1989 SCF for years prior to 1989 and filling in missing values thereafter through interpolation. The estimated marginal propensity to consume is lower than the baseline estimate but still relatively high.

By basing equity shares on a micro data set like the SCF instead of aggregate data, we can also consider the effect of heterogeneity in portfolios. A large literature suggests that portfolio composition differs systematically across households with different income and wealth (e.g., Friend and Blume, 1975, and Bertaut and Starr-McCluer, 2000), and the lower panel of figure 4 shows that equity shares in the SCF tend to fall as wealth rises.²¹ For the third column of table 7, we recalculate capital gains using quartile-specific measures of the equity share from the SCF. Allowing for this type of heterogeneity in portfolio composition raises the estimated marginal propensity to consume slightly relative to the results in the previous column.

The final column of table 7 explores the bias that could arise from other assets in reported securities having returns that are positively correlated with the stock market. Because bond returns had a positive correlation with stock returns during our sample, we may have understated household wealth changes measuring only the gains in stocks. Here we use a measure of capital gains that assumes an equity share of 100 percent, the equivalent to the extreme assumption that the returns on other assets have a perfect correlation with stock returns. The estimated long-run marginal propensity to consume drops to 0.087, and remains statistically significant at better than the 1 percent level.

Table 8 explores yet another possible source of bias in the results—the fact that many CE households may be experiencing (and reacting to) capital gains on assets held outside of reported “securities.” In particular, many CE households may be holding equities through defined contribution (DC) pension plans such as IRA’s or 401(k)’s. If DC pension holdings are positively correlated with other equity holdings, our measure of capital gains would understate the true wealth change associated with stock market movements and, in turn, lead to overstated

²¹ The narrowing of this differential in 1998 may be due to an especially high return on equities held by richer households during the previous several years. Of course, the equity share of total securities is just one way in which portfolio composition may differ with wealth. Our results could also be biased if the specific equities held by rich households tend to have higher risk and expected returns than those held by lower-wealth households. Unfortunately, the SCF does not provide the level of detail necessary to explore this issue.

estimates of the marginal propensity to consume.²² Unfortunately, the CE provides essentially no information about participants’ pensions. Aggregate data is of little use as well, because the aggregate role of DC holdings heavily reflects the importance of these assets for the wealthiest households. Thus, we again turn to the SCF, and calculate the average of:

$$\frac{\text{stocks in reported securities} + \text{stocks in DC pension plans}}{\text{stocks in reported securities}} \quad (13)$$

for a sample of households restricted comparably to our CE sample. The 1989 SCF was the first wave with enough detail to do this calculation, so we estimate values for 1983 through 1988 by assuming this ratio increased as much as the corresponding ratio for total DC holdings and total securities over that period. For missing years after 1989, we again used interpolation.

The second, third, and fourth columns of Table 8 show results when our previous measures of capital gain are scaled up for DC pensions according to equation (13). As in table 7, we repeat the exercise using the aggregate equity share, the share from the SCF, and a share of 100 percent.²³ In all three cases, the estimated marginal propensity to consume is much lower than in the previous tables and still statistically significant at the five percent level. Since the final column involves an extreme assumption about how closely non-equity returns are correlated with the stock market, the estimated long-run marginal propensity to consume—at just over 5 cents—puts a lower bound on the consumption response for this sample.

Heterogeneous Households

Tables 6 through 8 present a range of estimates of the long-run marginal propensity to consume that exceed the 0.03 to 0.05 range commonly estimated from aggregate time series data. This may imply that the marginal propensity to consume of the wealthy households left out of our sample is lower than the aggregate value — in other words, that marginal propensities to

²² In contrast with this exercise, Thaler (1990) argues that consumption might not be very responsive to capital gains on DC pensions because households engage in “mental accounting.”

²³ In the latter case the scale factor is based on total DC holdings and total securities.

consume fall with wealth.²⁴ One possible explanation is that uncertainty creates a consumption function that is convex with respect to wealth, as discussed earlier. Alternatively, returning to equation (3), lower-wealth households may simply be more impatient, or (perhaps by putting a lower value of the utility of future generations) have a shorter time horizon.

We can do back-of-the-envelope calculations to determine the marginal propensities to consume of the very wealthy households excluded from our sample ($mpc^{sec \geq \$100K}$) that would be consistent with our estimates for other households ($mpc^{sec < \$100K}$) and aggregate estimates (mpc^{agg}). On average, in the 1983, 1989, 1992, 1995, and 1998 SCF's, households with securities greater than \$100,000 held 84 percent of equities. Assuming that their stock returns are the same as those for other households, the marginal propensity to consume of wealthy households is:

$$mpc^{sec \geq \$100K} = \frac{1}{0.84} (mpc^{agg} - (1 - 0.84) mpc^{sec < \$100K}) . \quad (14)$$

Table 9 show the implicit value of $mpc^{sec \geq \$100K}$ under different assumptions about the aggregate marginal propensity to consume and the marginal propensity to consume for our sample. The estimates suggest that a \$1 gain in stock market wealth leads very wealthy households to raise their consumption between 1 cent and 5 cents.

6. Conclusions

Our results imply the aggregate relationship between consumption and stock market wealth is consistent with a "direct" view of wealth effects, in which changes in total consumption stem from changes in the consumption of households that own stocks. The consumption growth of *Consumer Expenditure Survey* households holding securities has a strong positive correlation with both contemporaneous and lagged movements in stock prices. The effect appears to taper off after a couple of years, much the same as in aggregate data. Meanwhile, our findings do *not* support an important role for "indirect" wealth effects: the consumption growth of households

²⁴Dynan, Skinner, and Zeldes (2000) show that *average* propensities to consume decline with permanent income and explore a variety of explanations for this finding.

that do not hold securities has little correlation with movements in stock prices.

Distinguishing between the direct and indirect channels is important for several reasons beyond the basic goal of better understanding household behavior. If wealth is not causal for consumption, then a decline in the stock market following a monetary tightening would best be interpreted as a symptom of an expected future slowdown in spending, rather than a cause. Further, the implications of a sharp correction in stock prices might differ if stock price changes matter primarily through the indirect channel because large changes in prices might cause disproportionate revisions in expectations of future economic conditions. In addition, the type of household behavior that drives aggregate wealth effects will determine how the aggregate marginal propensity to consume out of wealth might change over time in response to changes in the distribution of holdings or institutional features that make such holdings more or less liquid.

The paper also estimates the marginal propensity to consume out of stock market wealth by constructing household-specific measures of capital gains. For households with reported securities less than \$100,000 (a limitation imposed by the CE's topcoding practices), our estimates imply that a \$1 capital gain raises the level of consumption by between 5 and 15 cents, with the effect occurring gradually over a couple of years. All of our point estimates of the marginal propensity to consume are highly statistically significant. We are inclined to put more weight on the lower part of this range, as those estimates are based on capital gains measures adjusted for equities held through defined contribution pension plans, an important part of many households' portfolios. Still, these estimates exceed the 0.03 to 0.05 range typically derived from aggregate data, suggesting that the very wealthy households excluded from our sample may have lower marginal propensities to consume.

Because of the limitations of our data, our results shed light on only a portion of the household-level underpinnings of the effect of stock market movements on the macroeconomy. Obviously, one area where more work is needed is using household-level data to estimate the response of very wealthy households to capital gains: over time, the change in the topcoding practices of the CE in 1995 will facilitate such analysis, as more waves of the survey become available with observations on the stockholdings of the rich. Also, as our results pertain only to spending on nondurable goods and services, future research should be directed at other ways in

which households may respond to wealth gains, including raising their expenditures on durable goods and housing, and adjusting their labor supply.

References

- Attanasio, Orazio. 1994. "Personal Saving in the United States," in James M. Poterba, ed., *International Comparisons of Household Saving*. Chicago and London: University of Chicago Press, 1994: 57-123.
- Attanasio, Orazio, James Banks, and Sarah Tanner. 1998. "Asset Holdings and Consumption Volatility." National Bureau of Economic Research working paper no. 6567 (May).
- Bertaut, Carol and Martha Starr-McCluer. 2000. "Household Portfolios in the United States." Federal Reserve Board Finance and Economics Discussion Series working paper no. 2000-26 (April).
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy. 1999. "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." Rodney L. White Center for Financial Research working paper no. 23-99 (October).
- Brayton, Flint and Peter Tinsley. 1996. "A Guide to FRB/US." Federal Reserve Board Finance and Discussion Series working paper no. 1996-42 (October).
- Carroll, Christopher D. and Lawrence H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance*. Chicago and London: University of Chicago Press, 1991: 305-43.
- Carroll, Christopher D. and Miles S. Kimball. 1996. "On the Concavity of the Consumption Function." *Econometrica* 64, no. 4 (July): 981-92.
- Cheng, Ing-Haw and Eric French. 2000. "The Effect of the Run Up in the Stock Market on Labor Supply." *Economic Perspectives* (Fourth Quarter).
- Davis, Morris and Michael Palumbo. 2000. "A Primer on the Economics and Time-Series Econometrics of Wealth Effects: A Primer." Federal Reserve Board.
- Dynan, Karen E. 2000. "Habit Formation in Consumer Preferences: Evidence from Panel Data." *American Economic Review* 90, no. 3 (June): 391-406.
- Dynan, Karen E., Jonathan Skinner and Stephen P. Zeldes. 2000. "Do the Rich Save More?" National Bureau of Economic Research working paper no. 7906 (September).
- Friend, Irwin, and Marshall E. Blume. 1975. "The Demand for Risky Assets," *American Economic Review* 70, no. 5 (December): 900-22.
- Fuhrer, Jeffrey C. 2000. "Habit Formation in Consumption: Implications for Monetary-Policy Models." *American Economic Review* 90, no. 3 (June): 367-390.

- Gabaix, Xavier and David Laibson, "The 6D Bias and the Equity Premium Puzzle." Mimeo, Harvard University (October).
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *The Journal of Political Economy* 86, no. 6 (December): 971-987.
- Juster, F. Thomas, Joseph Lupton, James P. Smith and Frank Stafford. 1999. "Savings and Wealth; Then and Now." Mimeo, University of Michigan (October).
- Kennickell, Arthur B. and Martha Starr-McCluer. 1997. "Retrospective Reporting of Household Wealth: Evidence from the 1983-89 Survey of Consumer Finances." *Journal of Business and Economic Statistics* 15, no. 3 (October): 452-463.
- Kiley, Michael T. 2000. "Identifying the Effect of Stock Market Wealth on Consumption: Pitfalls and New Evidence." Mimeo, Federal Reserve Board (July).
- Lawrance, Emily C. 1991. "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99, no. 1 (February): 54-77.
- Ludvigson, Sydney and Charles Steindel. 1999. "How Important is the Stock Market Effect on Consumption?" *Federal Reserve Bank of New York Economic Policy Review* 5, no. 2 (July): 29-52.
- Maki, Dean M. 2001. "Household Debt and the Tax Reform Act of 1986." *American Economic Review* 91, no. 1: 305-19.
- Maki, Dean M. and Michael G. Palumbo. 2001. "Disentangling the Wealth Effect: A Cohort Analysis of Household Saving in the 1990s." Federal Reserve Board Finance and Discussion Series working paper no. 2001-21 (April).
- Mankiw, N. Gregory and Stephen P. Zeldes. 1991. "The Consumption of Stockholders and Nonstockholders." *Journal of Financial Economics* 29: 97-112.
- Otoo, Maria W. 1999. "Consumer Sentiment and the Stock Market." Federal Reserve Board Finance and Discussion Series working paper no. 1999-60.
- Parker, Jonathan A. 1999a. "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate" in Ben Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual, 1999*. MIT Press, Cambridge.
- Parker, Jonathan A. 1999b. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review* 89, no. 4 (September): 959-73.
- Parker, Jonathan A. 1999c. "The Consumption Function Re-Estimated." Mimeo, Princeton

University, August.

Parker, Robert and Clinton McCully. 2000. "The Role of Household Expenditure Surveys in Measuring Consumption." Presentation at conference on *Issues in Measuring Price Change and Consumption*, Bureau of Labor Statistics, June.

Poterba, James M. 2000. "Stock Market Wealth and Consumption." *Journal of Economic Perspectives* 14, no. 2 (Spring): 99-118.

Poterba, James and Andrew Samwick. 1995. "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption." *Brookings Papers on Economic Activity* no. 2: 295-357.

Sabelhaus, John. 1993. "What is the Distributional Burden of Taxing Consumption." *National Tax Journal* 46, no. 3 (September): 331-44.

Sabelhaus, John. 1998. "Using the Consumer Expenditure Survey for Research on Saving and Consumption." Mimeo, Congressional Budget Office (May).

Shleifer, Andrei. 1995. "Comments on Poterba and Samwick's 'Stock Ownership Patterns, Stock Market Fluctuations, and Consumption'." *Brookings Papers on Economic Activity* no. 2: 358-59.

Souleles, Nicholas S. 1999a. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89, no. 4 (September): 947-73.

Souleles, Nicholas S. 1999b. "Household Securities Purchases, Transactions Costs, and Hedging Motives." Rodney L. White Center for Financial Research working paper no. 24-99 (June).

Starr-McCluer, Martha. 2000. "Stock Market Wealth and Consumption." Mimeo, Federal Reserve Board (February).

Thaler, Richard H. 1990. "Anomalies: Saving, Fungibility, and Mental Accounts." *The Journal of Economic Perspectives* 4, no. 1. (Winter): pp. 193-205.

Vissing-Jørgensen, Annette. 1999. "Limited Stock Market Participation and the Equity Premium Puzzle. Mimeo, University of Chicago (December).

Zeldes, Stephen P. 1989a. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97, no. 2 (April): 305-46.

Zeldes, Stephen P. 1989b. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104, no. 2 (May): 275-98.

Table 1
Relationship between Consumption Growth and Current Stock Returns
for Stockholders and Other Households

	Full Sample	Stockholders (SH) defined as households with securities > \$0		Stockholders (SH) defined as households with securities > \$1000		Stockholders (SH) defined as households with securities > \$10,000	
		SH	other HHs	SH	other HHs	SH	other HHs
$\Delta \ln Wilshire_t$	-.035 (.030)	.124 (.069)	-.071 (.033)	.157 (.075)	-.072 (.032)	.181 (.100)	-.058 (.031)
<i>Age</i>	-.16 (.10)	-.03 (.26)	-.19 (.11)	-.01 (.29)	-.19 (.11)	.02 (.39)	-.20 (.10)
<i>Age</i> ² /1000	1.69 (.96)	-.27 (2.51)	2.12 (1.04)	-.51 (2.73)	2.13 (1.03)	-.38 (3.61)	2.04 (1.00)
<i>Income</i> /1000	-.004 (.011)	.019 (.021)	-.027 (.013)	.026 (.022)	-.025 (.013)	.012 (.028)	-.018 (.012)
<i>High School</i>	-.70 (.69)	2.39 (2.29)	-1.02 (.73)	1.89 (2.48)	-.88 (.72)	2.09 (3.18)	-.92 (.71)
<i>College</i>	.97 (.87)	1.34 (2.42)	1.13 (.98)	1.00 (2.62)	1.16 (.96)	2.02 (3.37)	.70 (.92)
$\Delta \ln Earnings$.020 (.003)	.005 (.008)	.022 (.004)	-.000 (.009)	.022 (.004)	-.006 (.012)	.022 (.003)
p-val. for F-test of month dummies	.000	.000	.000	.001	.000	.117	.000
p-val. for F-test of year dummies	.002	.517	.001	.616	.000	.412	.001
R ²	.009	.012	.010	.011	.010	.013	.010
Number of obs.	31,395	5,945	25,450	5,202	26,193	3,135	28,260

Note. Standard errors in parentheses. Dependent variable is the log difference of nondurables and services consumption as defined in the text. Regressions also include family size and dummy variable for a white head.

Table 2
Relationship between Consumption Growth and Current and Lagged Stock Returns
for Stockholders and Other Households

	Full Sample	Stockholders (SH) defined as households with securities > \$0		Stockholders (SH) defined as households with securities > \$1000		Stockholders (SH) defined as households with securities > \$10,000	
		SH	other HHs	SH	other HHs	SH	other HHs
$\Delta \ln Wilshire_t$.001 (.044)	.208 (.101)	-.050 (.049)	.241 (.109)	-.049 (.048)	.339 (.145)	-.038 (.046)
$\Delta \ln Wilshire_{t-9}$.079 (.051)	.212 (.118)	.046 (.057)	.215 (.127)	.049 (.056)	.326 (.171)	.049 (.054)
$\Delta \ln Wilshire_{t-18}$.045 (.047)	.134 (.109)	.022 (.052)	.113 (.118)	.028 (.051)	.239 (.160)	.020 (.049)
$\Delta \ln Wilshire_{t-27}$.025 (.037)	.109 (.086)	.009 (.041)	.085 (.093)	.014 (.041)	.128 (.124)	.012 (.039)
<i>Age</i>	-.16 (.10)	-.04 (.26)	-.19 (.11)	-.01 (.29)	-.19 (.11)	.00 (.39)	-.20 (.10)
<i>Age</i> ² /1000	1.69 (.96)	-.24 (2.51)	2.12 (1.04)	-.50 (2.73)	2.13 (1.03)	-.21 (3.61)	2.03 (1.00)
<i>Income</i> /1000	-.004 (.011)	.019 (.021)	-.027 (.013)	.026 (.022)	-.025 (.013)	.014 (.028)	-.018 (.012)
<i>High School</i>	-.69 (.69)	2.52 (2.29)	-1.01 (.73)	2.02 (2.48)	-.87 (.72)	2.30 (3.19)	-.91 (.71)
<i>College</i>	.98 (.87)	1.44 (2.42)	1.13 (.98)	1.10 (2.62)	1.16 (.96)	2.19 (3.37)	.70 (.92)
$\Delta \ln Earnings$.020 (.003)	.005 (.008)	.022 (.004)	-.000 (.009)	.022 (.004)	-.006 (.012)	.022 (.003)
p-val. for F-test of month dummies	.000	.000	.000	.000	.000	.097	.000
p-val. for F-test of year dummies	.005	.672	.002	.719	.001	.605	.001
R ²	.009	.012	.010	.012	.010	.014	.010
Number of obs.	31,395	5,945	25,450	5,202	26,193	3,135	28,260

Note. Standard errors in parentheses. Dependent variable is the log difference of nondurables and services consumption as defined in the text. Regressions also include family size and dummy variable for a white head.

Table 3
Relationship between Earnings and Hours Growth and Stock Returns
For Stockholders and Other Households

	Full Sample	Stockholders (SH) defined as households with securities > \$0		Stockholders (SH) defined as households with securities > \$1000		Stockholders (SH) defined as households with securities > \$10,000	
		SH	other HHs	SH	other HHs	SH	other HHs
----- <i>Dependent Variable: Earnings Growth</i> -----							
$\Delta \ln Wilshire_t$.041 (.075)	.037 (.161)	.044 (.084)	-.009 (.172)	.048 (.083)	-.293 (.223)	.076 (.079)
$\Delta \ln Wilshire_{t-9}$.075 (.087)	-.067 (.187)	.106 (.098)	-.133 (.199)	.105 (.097)	-.394 (.254)	.124 (.092)
$\Delta \ln Wilshire_{t-18}$.009 (.080)	-.035 (.173)	.015 (.090)	-.105 (.185)	.022 (.089)	-.357 (.245)	.048 (.085)
$\Delta \ln Wilshire_{t-27}$.084 (.063)	-.084 (.137)	.122 (.071)	-.090 (.146)	.116 (.070)	-.114 (.190)	.110 (.067)
Number of obs.	31,395	5,945	25,450	5,202	26,193	3,135	28,260
----- <i>Dependent Variable: Hours Growth</i> -----							
$\Delta \ln Wilshire_t$.005 (.054)	.175 (.110)	-.036 (.061)	.222 (.120)	-.040 (.060)	.214 (.168)	-.021 (.057)
$\Delta \ln Wilshire_{t-9}$	-.077 (.064)	.014 (.129)	-.100 (.073)	.070 (.140)	-.108 (.072)	.005 (.200)	-.089 (.068)
$\Delta \ln Wilshire_{t-18}$	-.078 (.059)	-.057 (.119)	-.084 (.067)	.011 (.131)	-.097 (.066)	-.157 (.187)	-.073 (.062)
$\Delta \ln Wilshire_{t-27}$	-.057 (.046)	-.083 (.094)	-.052 (.052)	-.044 (.102)	-.060 (.051)	-.250 (.143)	-.036 (.048)
Number of obs.	28,672	5,509	23,163	4,798	23,874	2,847	25,825

Note. Standard errors in parentheses. Regressions also include family size and dummy variable for a white head.

Table 4
Comparison of Consumer Expenditure Survey Data and NIPA Data

Year	CE after-tax income		CE total expenditures		CE “saving rate”	NIPA saving rate
	Billions	Percent of NIPA after-tax income	Billions	Percent of NIPA Outlays		
1993	\$3191	64.7	\$3071	67.0	3.8	7.1
1994	\$3383	65.5	\$3243	66.9	4.1	6.1
1995	\$3492	64.4	\$3327	65.0	4.7	5.6
1996	\$3633	64.0	\$3522	65.2	3.1	4.8
1997	\$3873	64.9	\$3676	64.3	5.1	4.2
1998	\$4111	65.1	\$3809	62.9	7.4	4.2
1999	\$4378	66.0	\$4016	61.9	8.3	2.2

Source. CE data from <ftp://ftp.bls.gov/pub/special.requests/ce/standard/y9399/multiyr.txt>.
NIPA data from Commerce Department.

Table 5
Shares of Households Holding Different Groupings of Securites

Year	CE	Survey of Consumer Finances	
	“Securities”	Stocks, bonds, and mutual funds	Stocks, bonds, mutual funds, and IRA’s/thrifts
1983	.21	.21	.36
1984	.20
1985	.21
1986	.21
1987	.22
1988	.22
1989	.23	.22	.43
1990	.20
1991	.20
1992	.21	.23	.44
1993	.21
1994	.22
1995	.21	.24	.49
1996	.23
1997	.25
1998	.27	.31	.56
1999	.26

Note. Results based on all households in each sample. SCF calculations were done with weights. “...” signifies that SCF wave not available for that year.

Table 6
Relationship Between Consumption Growth and Capital Gains

Controls:				
month and year dummies	x	x	x	x
tasteshifters	x	x	x	x
socioeconomic variables	x	x		x
Δ earnings	x			x
year x education dummies				x
<i>CAPGN</i> _{<i>it</i>}	.091 (.034)	.094 (.034)	.092 (.034)	.095 (.034)
<i>CAPGN</i> _{<i>it-9</i>}	.050 (.027)	.057 (.027)	.054 (.027)	.047 (.027)
<i>CAPGN</i> _{<i>it-18</i>}	.033 (.027)	.036 (.027)	.033 (.027)	.030 (.027)
<i>CAPGN</i> _{<i>it-27</i>}	-.043 (.036)	-.039 (.036)	-.038 (.036)	-.040 (.036)
Sum of <i>CAPGN</i> coefficients	.132 (.042)	.148 (.042)	.142 (.041)	.132 (.042)
R ²	.012	.009	.008	.013
Number of obs.	33,110	33,110	33,110	33,110

Note. Standard errors in parentheses. Dependent variable is the change in nondurables and services consumption as described in text.

Table 7
Relationship Between Consumption Growth and Capital Gains
Under Different Assumptions about the Equity Share

	Baseline* (Aggregate equity share)	Equity share from SCF		Assume equity share equals 100 percent
		Same share for all households	Shares vary by securities quartile	
$CAPGN_{it}$.091 (.034)	.076 (.027)	.088 (.031)	.059 (.021)
$CAPGN_{it-9}$.050 (.027)	.044 (.022)	.050 (.025)	.034 (.017)
$CAPGN_{it-18}$.033 (.027)	.026 (.022)	.030 (.025)	.020 (.016)
$CAPGN_{it-27}$	-.043 (.036)	-.034 (.028)	-.041 (.033)	-.027 (.021)
Sum of $CAPGN$ coefficients	.132 (.042)	.112 (.035)	.127 (.040)	.087 (.001)
R ²	.012	.012	.012	.012
Number of obs.	33,110	33,110	33,110	33,110

* Same results as in the first column of table 6.

Note. Standard errors in parentheses. Dependent variable is the change in nondurables and services consumption as described in text. Regressions also include month and year dummies, tasteshifters, socioeconomic variables, and change in earnings.

Table 8
Relationship Between Consumption Growth and Capital Gains
Scaling up to Reflect Holdings in Defined Contribution Pensions

	Baseline* (no scaling for DC holdings)	Scaling up using information from SCF		
		Aggregate equity share	Equity share from SCF	Assume equity share equals 100 percent
<i>CAPGN</i> _{<i>it</i>}	.091 (.034)	.110 (.035)	.094 (.029)	.048 (.015)
<i>CAPGN</i> _{<i>it-9</i>}	.050 (.027)	.048 (.033)	.044 (.027)	.027 (.014)
<i>CAPGN</i> _{<i>it-18</i>}	.033 (.027)	-.012 (.035)	-.010 (.029)	.004 (.014)
<i>CAPGN</i> _{<i>it-27</i>}	-.043 (.036)	-.073 (.043)	-.059 (.035)	-.027 (.017)
Sum of <i>CAPGN</i> coefficients	.132 (.042)	.073 (.038)	.070 (.033)	.052 (.019)
R ²	.012	.012	.012	.012
Number of obs.	33,110	33,110	33,110	33,110

* Same results as in the first column of table 6.

Note. Dependent variable is the change in the baseline measure of nondurables and services consumption. Standard errors in parentheses. Sample selection and variable definitions described in the text. Regressions also include month and year dummies, tasteshifters, socioeconomic variables, and change in earnings.

Table 9
Implied MPC of Households with Securities Greater than \$100,000
Under Different Assumptions about Lower Wealth Households and the Aggregate MPC

	Aggregate MPC		
	.03	.04	.05
MPC of Lower Wealth Households			
.06	.024	.036	.048
.10	.017	.029	.040
.14	.009	.021	.033

Figure 1

Saving Rate and Wealth-to-Income Ratio

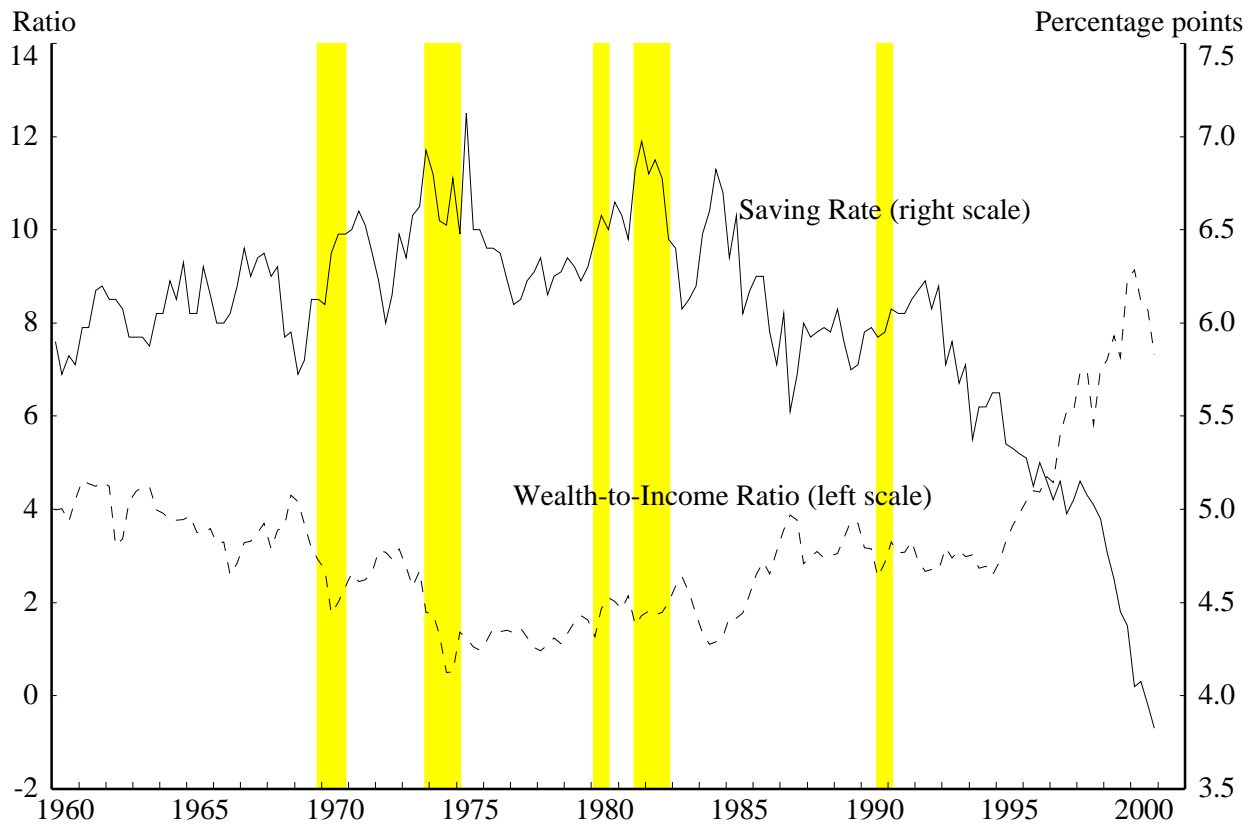


Figure 2
Timing of Consumption Changes and Wilshire Changes




Figure 3
Timing of Consumption Changes and Capital Gains




Figure 4

