Consumption-Wealth Comovement of the Wrong Sign

by

James J. Choi
Harvard University

David Laibson
Harvard University and NBER

Brigitte C. Madrian
University of Pennsylvania and NBER

Andrew Metrick
University of Pennsylvania and NBER

Date: September 9, 2003

We thank Hewitt Associates for their help in providing the data. We are particularly grateful to Lori Lucas, Jim McGhee, Scott Peterson, and Yan Xu, some of our many contacts at Hewitt. Outside of Hewitt, we have benefited from conversations with Thomas Knox and Bill Zame. We are also indebted to John Beshears, Fuad Faridi, John Friedman, Parag Pathak, Nelson Uhan, Kenneth Weinstein, and Lucia Wu for their excellent research assistance. Financial support from the National Institute on Aging (grants R01-AG021640 and R01-AG16605) is gratefully acknowledged. Choi acknowledges financial support from a National Science Foundation Graduate Research Fellowship and National Institute on Aging Grant T32-AG00186. Laibson acknowledges financial support from the Sloan Foundation.

Consumption-Wealth Comovement of the Wrong Sign

Abstract: Economic theory predicts that an unexpected wealth windfall should increase consumption as soon as the windfall is received. We test this prediction by using administrative records on over 40,000 401(k) accounts. Contrary to theory, we estimate a negative short-run marginal propensity to consume out of orthogonal 401(k) capital gains shocks. Our findings suggest that many investors are influenced by a reinforcement learning heuristic that causes high returns to encourage saving and low returns to discourage saving. These results help explain why consumption covariance with equity returns is so low, giving rise to the equity premium puzzle.

James J. Choi
Department of Economics
Harvard University
Littauer Center
Cambridge, MA 02138
james_choi@post.harvard.edu

David Laibson
Department of Economics
Harvard University
Littauer Center
Cambridge, MA 02138
dlaibson@harvard.edu

Brigitte C. Madrian
Department of Business and Public Policy
University of Pennsylvania, Wharton School
3620 Locust Walk
Philadelphia, PA 19104
bmadrian@wharton.upenn.edu

Andrew Metrick
Department of Finance 2300 SH-DH
University of Pennsylvania, Wharton School
3620 Locust Walk
Philadelphia, PA 19104
metrick@wharton.upenn.edu

Economic theory predicts that an unexpected wealth windfall should increase consumption as soon as the windfall is received. In this paper, we present empirical evidence that contradicts this prediction in an important domain: 401(k) accounts, which represented $1.75 trillion in assets held by 45 million workers at year-end 2001 (Holden and Vanderhei, 2003). Using administrative data on over 40,000 401(k) accounts across five companies, we measure cross-sectional variation in individual consumption growth and unexpected capital gains. The size of the capital gain in one's own portfolio relative to another investor's portfolio is irrelevant for forecasting future returns. Therefore, only an income effect should be identified from our regressions. Contrary to theory, the resulting estimates of the short-run marginal propensity to consume (MPC) out of these orthogonal capital gains shocks are negative. These results are consistent with Starr-McCluer (2002), who finds that 11.6% of stockholders report that the 1990s bull market caused them to save more, while only 3.4% say they saved less.

We interpret our negative MPC estimates as the consequence of a reinforcement learning mechanism. Reinforcement learning models have had success in predicting the pattern of play in experimental settings (Erev and Roth 1998, Charness and Levin 2003). In our context, reinforcement encourages saving in response to high returns and discourages saving in response to low returns, working in opposition to the classical income effect. Since most of our wealth shocks are driven by equities, our results help explain why consumption covariance with equity returns is so low, giving rise to the equity premium puzzle.

The outline of the paper is as follows. Section I describes our data. Section II lays out our empirical methodology. Section III presents our main results and reconciles them with previous literature that has found positive MPCs out of windfalls. Section IV considers two broad classes of objections to our results. The first is that people adjust their consumption along margins that our consumption measure does not capture. The second is that our results are driven by wealth shocks outside of the 401(k) which we do not observe. We find no evidence to support these objections. Section V concludes.

I. Data description

Our data come from a large benefits administration and consulting firm. We have panel data for five companies that start when our data provider became the plan administrator at each firm and end at year-end 2000. These data contain the date, amount, and type of every transaction made in the 401(k) plans by every participant. In addition, we have cross-sectional snapshots of age, annual salary, date of hire, gender, marital status, 401(k) asset allocations, and elected 401(k) contribution rates as of year-ends 1998, 1999, and 2000 for those actively employed at the companies on those dates.

Table 1 gives summary statistics as of year-end 2000 for our companies, which we code-name Company A through E. Our sample consists of large firms that span a wide range of industries. The employees are on average 42.9 years old and earn $55,292 a year. By comparison, the March 2001 Current Population Survey reports an average age of 40.8 and salary of $45,656 among full-time workers whose company employs over 1,000 workers and offers some kind of retirement plan. The average 401(k) participation rate across the firms is 79%, which is close to the 2000 national participation rate of 80% found by the Profit Sharing/401(k) Council of America (2001), and the average balance of participants is $65,964, which is similar to Holden and Vanderhei's (2001) reported average year-end 2000 balance of $61,207 among plans with more than 10,000 participants.

All of our companies offer matching contributions ranging from 25 cents to a full dollar for each dollar contributed to the 401(k) by the employee up to a threshold, although Company C did not introduce its match until 2000. The early withdrawal and loan provisions of our plans are generous by industry standards, so participants' 401(k) balances are relatively liquid.[1] Finally, all of the plans allow changes to the elected contribution rate and asset allocation on a daily basis. Changes can be made 24 hours a day using the phone or Internet. Therefore, the direct transaction costs involved in changing one's savings rate in these plans are minimal.

II. Empirical methodology

Our objective is to estimate the relationship between an orthogonal wealth shock in year t and consumption growth between year $t - 1$ and year t. The key assumption is that consumption adjustments are observable through changes in the 401(k) contribution rate. This assumption is plausible because most households should be doing all of their saving in the 401(k), since its tax benefits and employer match (for the plans in our data) make it the most attractive savings

[1] The U.S. Department of Labor (2003) reports that in 2000, 40% of full-time employees with savings and thrift plans in private industry were not allowed to take early in-service withdrawals for any reason, and an additional 29% could only take hardship withdrawals. The Profit Sharing/401(k) Council of America (2001) reports that 14% of plans did not permit loans in 2000.

vehicle available to them. Consistent with this assumption, among 401(k)-holding households earning less than $70,000 a year in the 2001 Survey of Consumer Finances—a sample roughly comparable to the one we will use in our analysis—the median household has less than one month's income in net financial assets outside the 401(k).[2] It is only at the 80th percentile that households have one year's income outside the 401(k), and this probably overstates outside asset holdings in our sample; because of the generosity of our 401(k) plans' early withdrawal and loan provisions, the need to maintain a precautionary stock of wealth outside the 401(k) is substantially mitigated. We will explore the possibility that the 401(k) contribution rate is not the savings adjustment margin in section IV.

Normalizing by income, our reduced-form expression for an individual's consumption growth from year $t-1$ to year t is

$$\frac{C_t - C_{t-1}}{Y_{t-1}} = \beta_1 \frac{Shock_t}{Y_{t-1}} + \beta_2 \frac{Shock_{t-1}}{Y_{t-1}}, \tag{1}$$

where C_t is annualized consumption flow during the last pay cycle[3] of year t, Y_{t-1} is annualized salary flow during the last pay cycle of year $t-1$, and $Shock_t$ is the wealth shock in dollars accrued from the beginning to the end of year t. The main coefficient of interest is β_1. If we use the simplification that a wealth shock changes consumption by a constant fraction of that shock each subsequent period, then we can interpret β_1 as the one-year marginal propensity to consume out of wealth. β_2 should equal zero if consumption responds to wealth innovations without a lag, as is commonly assumed in theory.

We will now transform this expression in order to make the dependent variable identifiable from our data. Let S_t be the 401(k) savings/contribution rate. We assume that

$$C_t = Y_t(1 - S_t) + k, \tag{2}$$

where k is a constant that represents consumption funded by income earned outside the company or savings outside the 401(k). Substituting this expression into (1) and simplifying yields

[2] On the assets side, we include CDs, bonds, savings bonds, publicly traded stock, mutual funds, cash value life insurance, other managed accounts, transactions accounts, non-401(k) pension accounts, and miscellaneous assets. For liabilities, we subtract credit card debt, non-home-equity lines of credit, business loans, education loans, other consumer loans, margin loans, loans against life insurance policies, loans against non-401(k) pension accounts, and car loans.

[3] In our sample, a pay cycle is typically two weeks.

$$S_{t-1} - S_t + \left(\frac{Y_t}{Y_{t-1}} - 1 \right)(1 - S_t) = \beta_1 \frac{Shock_t}{Y_{t-1}} + \beta_2 \frac{Shock_{t-1}}{Y_{t-1}}. \qquad (3)$$

We add the before-tax and (if the plan offers the option) after-tax 401(k) contribution rates in effect for the last pay cycles of 1998, 1999, or 2000 to calculate S_t. Note that the resulting consumption measure spans only the last two weeks of each year and provides a maximum of three consumption observations per individual, which translates into two consumption growth observations. Y_t is set to annualized salary at the end of year t in 1998 dollars, assuming that the salary was paid in nominally equal amounts each pay cycle and deflating by the monthly CPI series.

In order to identify wealth shocks, $Shock_t$ and $Shock_{t-1}$, we cannot simply measure the change in 401(k) balances between year-ends. There are two reasons. First, contributions to the 401(k) plan should not be counted as increases in wealth, since they are merely a transfer from human wealth to financial wealth. Second, even if we were to isolate capital gains from contributions, a portion of those capital gains is expected and thus should not affect consumption growth. We must therefore subtract off the expected component of capital gains.[4] With these considerations in mind, we define the wealth shock accrued over a year as the difference between the realized and expected dollar capital gain in each asset, summed over all assets.

Thus, to compute the wealth shock we must measure both realized and expected capital gains. Capital gains are calculated in the standard manner: for each asset, we multiply monthly percent returns by the participant's dollar holdings at the end of the prior month. Assets are defined at the level of individual mutual funds and employer stock holdings. Dollar holdings are deflated by the same monthly CPI series used to deflate income, Y_t. The sum of the monthly dollar returns for all assets and all months yields the capital gain for the year.

We estimate expected percent returns using two different measures. The first measure, the "typical return," is defined for three asset classes—equities (including employer stock), bonds, and cash. The typical return for an asset class in a particular company during a particular month is the average return that 401(k) participants of that company realized in that asset class during that month.[5,6] The typical return is motivated by the intuition that the deviation of a

---

[4] We have also estimated the regressions simply using total capital gains as the wealth shock measure, and our results are qualitatively unchanged.

[5] The average is taken over participants with holdings in the relevant asset class.

mutual fund's return from the contemporaneous average return in its asset class—the "typical" return—is unforecastable. Therefore, the typical return wealth shock is the difference between what the participant actually gained from holding a fund during the month and what he would have realized had he instead held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year.

We call our second expected-return measure the "adaptive return" because it assumes that the expected return for a mutual fund is equal to its asset class's lagged ten-year return.[7] Benartzi (2001) presents survey evidence that expectations of future employer stock returns are highly correlated with lagged employer stock returns. In accordance with these responses, Benartzi (2001) and Choi, Laibson, Madrian, and Metrick (2003) find that the fraction of 401(k) contributions going to employer stock is highly correlated with the stock's lagged performance. Anecdotally, it seems that the bear market of 2000 was a surprise to many people because the bull market of the preceding decade had created expectations of many additional years of high stock market returns. Table 2 gives the indexes used to measure the eleven lagged asset class returns. Note that we have much finer asset class categories for this measure than for typical returns.

In general, consumption growth is expected to be non-zero even in the absence of a wealth shock because of intertemporal substitution motives. Furthermore, the utility function shifts predictably with demographic variables, causing changes in consumption expenditure. If our wealth shock measures were orthogonal to all information useful for predicting utility function shifts, expected asset returns, and the elasticity of intertemporal substitution, we could ignore those factors in our regressions. We believe that our typical return wealth shock comes close to meeting this criterion in population. However, certain funds idiosyncratically outperform or underperform in sample, so our wealth shock measures may be correlated with other variables that predict both consumption growth and fund choices. In light of these considerations, we modify (3) to arrive at our final regression equation,

[6] If the plan offers only one fund in a particular asset class—not an unusual situation for cash and fixed-income funds—then expected returns and realized returns would always be the same in that asset class and the wealth shock arising from holding this asset class would always be zero.
[7] We have also run our regressions using one-year lagged returns, and our results are qualitatively unchanged.

$$S_{t-1} - S_t + \left(\frac{Y_t}{Y_{t-1}} - 1 \right) \left(1 - S_t \right) = \alpha + \beta_1 \frac{Shock_t}{Y_{t-1}} + \beta_2 \frac{Shock_{t-1}}{Y_{t-1}} + \boldsymbol{\gamma}' \mathbf{X}_{t-1} + \varepsilon_t . \tag{4}$$

We estimate this equation stacking our data across time periods and using a tobit to account for the fact that 401(k) contribution rates cannot be negative or above the plan maximum. If S_t is the plan's maximum contribution rate or S_{t-1} is zero, the observation is considered left-censored; if S_t is zero or S_{t-1} is the maximum contribution rate, the observation is considered right-censored. We also estimate (4) allowing the intercept term α to vary by individual using an OLS difference estimator.

In order to be in the tobit regressions, an employee must have been actively employed at the firm and enrolled in the 401(k) plan for at least two consecutive complete calendar years between 1998 and 2000. We include employees whose contribution rate or plan balances are zero, provided that he or she had positive balances at some time in the past. We require that individuals have salaries greater than \$20,000 in the prior year because a large fraction of those with salaries under \$20,000 are part-time employees. We also trim observations where one-year income growth is above 30% or below –20%, which roughly corresponds to removing the top 2% and bottom 2% of the income growth distribution.[8] These deleted outliers are usually caused by changes in labor force participation that did not entail complete separations from the firm. Finally, we drop individuals if their salary is high enough in the prior year that they could exceed the statutory annual dollar limit on 401(k) contributions (\$10,000 in 1998 and 1999) by contributing at the plan's maximum contribution rate. The reason we impose this selection rule can be illustrated by the following example. Suppose a highly-paid employee contributes enough that he hits the dollar limit midway through the year. Assuming that he continues to save a positive amount after hitting the limit, these additional savings enter an account outside the 401(k). If he decides to increase his consumption in December, this will likely be financed through a drop in his contributions to the outside account. This consumption flow increase will not be reflected in his 401(k) contribution rate.

These criteria leave us with a final tobit sample of 69,676 observations on 42,396 employees. The typical returns for the tobit regressions are estimated averaging over this sample.

[8] We have also run our regressions trimming the top 1% and bottom 1% of income growth, as well as the top 5% and bottom 5% of income growth. Those results are qualitatively similar to the ones presented in the tables.

We now discuss analogous sample selection criteria for our OLS difference regressions, which incorporate employee-level fixed effects. For these fixed-effect regressions we require that 1998 salary be above \$20,000 (thereby eliminating part-time workers as explained above). We also require that the employee be actively employed and enrolled in the 401(k) plan from January 1, 1998 to December 31, 2000. The cutoffs we use to trim salary growth outliers remain the same as in the original tobit sample. This leaves us with 54,560 observations on 27,280 employees in the fixed-effects sample. We estimate a separate set of typical returns for the fixed-effects regressions using this sample.

It turns out that including the vector of control variables, $\mathbf{X}_{t-1}$, does not qualitatively change the MPC estimates, but we describe these variables for completeness. For the tobit regressions, we include age, age squared, log of salary in 1998 dollars, and log of tenure at the company, all as of $t-1$. Additionally, we include company-year dummies. When running the employee fixed-effects regressions, we must change the $\mathbf{X}_{t-1}$ vector. First, we must eliminate age because it increases linearly and is thus not separately identified. Second, we interact the company dummies with only one year dummy.

III. Main results

A. Regression estimates

Table 3 reports summary statistics for our consumption growth and normalized wealth shock measures in the tobit sample. Median consumption growth is 1.0% of income, and the distribution is skewed right. Both normalized wealth shock measures have a median close to zero and a relatively narrow distribution. The 10th percentile of the normalized typical return wealth shock is a loss of 12.2% of a year's income, and the 90th percentile is a gain of 9.2% of a year's income. The spread is wider for the normalized adaptive wealth shock; the 10th and 90th percentiles are separated by 53.5% of annual income. We will show that this amount of variation in wealth shocks is sufficient to produce precise estimates of the MPC.

Table 4 presents the coefficients from estimating equation (4). We find that all four MPC estimates are negative; three of them are statistically significant at conventional levels and the fourth barely misses statistical significance with a p-value of 0.056. The point estimates range from –0.45% to –1.65%, which means that a positive wealth shock equal to one year of income will contemporaneously decrease annualized consumption flow by 0.45% to 1.65% of yearly income. There is no consistent evidence that consumption increases with a lag in response to positive wealth shocks. Three of the four lagged response estimates are negative, and the only significantly positive estimate of 1.04% is counterbalanced by significantly negative estimates of –0.41% and –1.69%.

A simple permanent income hypothesis model predicts that the MPC is higher (more positive) for the old than the young, since the young have a longer remaining lifespan over which to spread consumption of a windfall. We test whether this stylized fact holds in the data by estimating separate MPCs for those who are under 30 years old, between 30 and 39, between 40 and 49, between 50 and 59, and 60 or above, while constraining the other regression coefficients to be equal across age groups. The results are in Table 5, where we have omitted coefficient estimates for the non-shock variables because they are similar to those in Table 4.

We find that the MPC does increase almost monotonically with age, but even the oldest participants do not exhibit a positive MPC. For example, the first column of Table 5 reports the results for a tobit regression using the typical return wealth shock as the wealth shock variable. For this model, we estimate an MPC of –2.90% for employees under 30, and an MPC of –0.17% for employees over 60. The corresponding employee fixed effects regression yields an estimate of –10.25% for those under 30 and –0.79% for those over 60. The lagged MPC does not exhibit any age-based pattern for the tobit regressions, but it does increase close to monotonically with age when estimated using employee fixed effects.

Overall, there is no compelling evidence that anomalous MPC behavior is restricted to a particular subset of the population.[9] Even older employees have negative MPC point estimates in our regressions. However, it does appear that the negative MPCs are of greater magnitude for the young, a point that we will interpret below.

[9] We have also run regressions interacting $Shock_t$ and $Shock_{t-1}$ with both age and salary on the theory that higher-paid workers are more sophisticated and therefore more likely to conform to neoclassical predictions. Contrary to expectations, salary has a *negative* and significant effect on contemporaneous MPC, while age continues to have a positive and significant effect. There is no consistent pattern for the interaction of age and salary with $Shock_{t-1}$.

B. Interpretation and Reconciliation with Past MPC Estimates

Our estimates are in sharp contrast to past research that has found positive MPCs between 16% and 97% out of orthogonal cash windfalls (Bodkin 1959, Kreinin 1961, Landsberger 1966, Imbens et al. 2001). Ludvigson and Steindel (1999) focus on aggregate stock market movements and estimate a positive 3-4% MPC.

We square our results with these previous findings by interpreting negative MPCs as the consequence of a reinforcement learning mechanism. Reinforcement learning models posit that agents are more likely to repeat actions that have previously generated a high payoff. For example, if an investor experiences a high return in one asset class, the investor will reallocate money to that asset class. Financial analysts commonly refer to such reallocation as return-chasing.

Return-chasing can be generalized to cover savings and consumption choices. Specifically, if a consumer experiences a gratifyingly high return from her savings activity, she will allocate more resources to savings and less to consumption. This positive savings effect is offset by the standard income effect, which pushes the investor to cut her savings. In our data, the reinforcement effect appears to dominate, even though the returns from which the MPC is identified convey no information about future returns. In particular, the young—whose income effect is the weakest and whose stock of past reinforcements is the smallest—are the most swayed by the reinforcement effect, producing negative MPCs with the largest magnitude.

This theory reconciles all of the evidence on MPCs. In the case of war veteran payments and Holocaust reparations (Bodkin 1959, Kreinin 1961, Landsberger 1966), there was no direct financial investment that causally preceded the windfalls, so there was no reinforcement to invest more. For lottery winners (Imbens et al. 2001), the cash investment in lottery tickets is so minimal relative to the payout that the income effect prevails. Ludvigson and Steindel (1999) use aggregate stock market and consumption data, which are dominated (on a dollar-weighted basis) by investors who are typically more sophisticated then the unsophisticated 401(k) investors in our sample. Unsophisticated investors are particularly likely to follow the reinforcement learning heuristic, since they don't understand that return chasing does not work in financial markets, where past returns by and large do not predict future returns. By contrast, more sophisticated investors have a positive MPC out of orthogonal capital gains shocks. Consistent with this hypothesis, Starr-McCluer (2002) finds that households with more than $250,000 in stockholdings were much more likely to report that the 1990s bull market caused them to increase their spending.

IV. Robustness checks

In this section, we consider possible objections to our results. Subsections IV.A through IV.D test the possibility that the 401(k) contribution rate is not the relevant consumption adjustment margin for the people in our data. Subsection IV.E discusses and tests the effect of non-401(k) wealth shocks on our results. Subsection IV.F considers the story that consumption adjusts not through expenditures but through leisure. We find no evidence that weakens the force of our main result. Of the 28 MPC estimates presented in this section, 24 are negative, and 17 of these are statistically significantly. None of the four positive point estimates are statistically significant, and they average an economically negligible 0.06%.

A. The 401(k) is not the relevant margin for consumption adjustment

Our MPC estimate hinges on the assumption that the 401(k) contribution rate is the margin at which participants adjust their savings rate. If participants find it worthwhile to adjust savings through contributions to other asset accounts, then the 401(k) contribution rate response to retirement wealth shocks may be offset by activity elsewhere.

We test this alternative explanation by restricting our tobit sample to participants who at year-end $t - 1$ were contributing less than the threshold to which their employer would provide matching contributions. Analogously, we restrict our employee fixed-effects sample to those who were contributing less than the match threshold at year-end 1998. These participants face instantaneous marginal returns to saving in their 401(k) of 25% to 100%. It is difficult to imagine that there are alternative investment vehicles that offer comparable returns. Therefore, these employees have especially strong incentives to adjust their consumption expenditures exclusively through their 401(k) contribution rate.[10] Because Company C did not have a match until 2000, its participants are excluded from this analysis.

---

[10] One might be concerned that the bulk of the participants who are contributing less than the match threshold are new employees whose employer matches will not vest for a long time. If such employees have a high probability of leaving the firm before their vesting begins, this restricted sample would not face significantly higher marginal incentives to save in the 401(k). However, it turns out that only 7.9% of those in the restricted tobit sample are not vested at all at year-end 1998, and 81.2% are fully vested.

The results of this regression are found in Table 6. All four MPC estimates remain negative, and the two estimates using employee fixed effects are statistically significant and economically large at –2.45% and –1.43%. Among the lagged MPC estimates, half are negative and half are positive.

Participants may also respond to positive wealth shocks by purchasing a house. Because mortgage payments cover both the rental flow of housing services and the purchase of housing equity, the house becomes a savings vehicle. The amount of money that goes towards home equity would come out of the 401(k) contribution rate, thus creating a spurious negative MPC estimate. In order to check for this possibility, we restrict our tobit sample to those whose zip code didn't change between $t - 1$ and t, and we restrict our individual fixed effects sample to those whose zip code didn't change between year-end 1998 and year-end 2000. The results in Table 7 show that all four MPC estimates are negative and statistically significant. Three of the four lagged MPC estimates are negative, two of them significantly so.

B. Consumption adjustment is occurring through spouse's 401(k)

401(k) accounts may be the most attractive savings vehicles available to the employees in our data, but the 401(k) accounts we see in our data may not be the only ones available to them. The employee's spouse's 401(k) may be more attractive and hence the one that attracts the marginal dollar. We test this story in our tobit regressions by restricting the sample to participants who were unmarried at $t - 1$. For the individual fixed effects regressions, we restrict the sample to participants who were unmarried at year-end 1998. Because we do not have marital status data on employees at Company D, they are excluded from this analysis. The regression results are presented in Table 8. We find that all four MPC estimates are negative and statistically significant. Three of the four lagged MPC estimates are negative, and two of these negative estimates are statistically significant.

C. Consumption adjustment is occurring through in-service withdrawals

We have been identifying changes in consumption through changes in the contribution rate and salary. However, participants may be making their consumption expenditure adjustments through in-service withdrawals from their 401(k) instead. In practice, this is unlikely to be a significant factor, given that only 4.8% of our tobit sample made any in-service withdrawals from the beginning of 1998 to the end of 2000. The main complication with a withdrawals analysis is the difficulty of assigning a time period to the consumption stream. At one extreme, one can assume that withdrawals are rolled over into another account and not consumed until the future, in which case withdrawals don't matter at all for our analysis. At the other extreme, one can assume that the entire withdrawal is consumed immediately. An intermediate case is to assume that the withdrawal is annuitized and consumed slowly over time.

We run regressions with two different assumptions about the timing of withdrawal consumption. Because withdrawals are infrequent events, we do not use a two-week measure of consumption as we did in the main analysis. Instead, we use yearly consumption, defined as the year's total income plus withdrawals consumed minus 401(k) contributions. We assume either that all withdrawals net of rollovers into other accounts are consumed in the year of the withdrawal, or that 5% of a withdrawal net of rollovers is consumed each year starting in the year of the withdrawal.[11] Withdrawals before age 59½ are assessed a 10% early withdrawal penalty. The dependent variable in our regressions is the year-over-year change in consumption normalized by prior year salary. The explanatory variables remain the same as in the previous regressions, as do the income, income growth, and plan enrollment restrictions.

In the tobit regressions, we consider an observation left-censored if in year t, in-service withdrawals net of rollovers equals zero and the individual contributed the maximum allowable given his or her salary and plan contribution rate limits. In other words, the individual could not have saved any more in the 401(k) plan during the year. Conversely, an observation is considered right-censored if the above criteria are satisfied for year $t-1$. An observation is also considered right-censored if total balances in the plan at the end of year t plus rollovers in year t equals zero. That is, the individual could not have funded any more consumption from the 401(k). If these conditions are satisfied in the lagged year, the observation is considered left-censored.

Assuming that the entire net withdrawal is consumed in the year of the withdrawal generates implausibly large outliers in consumption growth. These are caused by large withdrawals that were not directly rolled over into another asset account by the 401(k)

[11] Our rollover measure is not comprehensive, however. We observe a withdrawal being rolled over into an IRA only when the employee asks that a check be sent directly from the employer to the IRA custodian. If an employee receives the withdrawal check him or herself and subsequently deposits some of the proceeds into an IRA, we do not observe this second transaction.

administrator. It is unlikely that such large sums were entirely consumed in one year. Therefore, we trim measured consumption growth in the top 1% and bottom 1%, which corresponds to constraining consumption growth to lie between –18.4% and 34.7%. We do not trim the dependent variable in the regressions that assume that the withdrawals are annuitized.

Table 9 presents the results when net withdrawals are assumed to be consumed immediately, and Table 10 presents the results when net withdrawals are assumed to be annuitized. Half the MPC estimates are negative, including the only statistically significant estimates, which average –1.67%. Among the positive estimates, the average magnitude is an economically trivial 0.06%. The lagged MPC estimates are split almost evenly between negative and positive numbers, and 4 out of the 6 significant estimates are negative. We conclude that accounting for in-service withdrawals cannot generate a positive MPC out of 401(k) wealth shocks.

D. Consumption adjustment is occurring through 401(k) loans

Withdrawals are not the only way to access money from one's 401(k) account; one can also take out loans against up to half the value of the 401(k). Loans, which can generally be taken for any reason, cause the balance in the 401(k) to fall by the amount of the principal, and then the employee must pay the loan back with interest to his or her own 401(k) account, typically through payroll deduction over a period of five to 25 years. Early repayment is possible with no penalties. Many companies charge loan origination fees, but none of our companies do. Unlike in-service withdrawals, loans are common in our sample: 29.8% of the tobit sample had a loan outstanding at some point between 1998 and 2000.

However, we do not believe that accounting for loans would capture significant consumption flow changes that are not already being measured by the contribution rate. Sundén and Surette (2000) find that only 8.5% of loans are used to finance non-durable consumption. 54.5% are used for durable expenditures on housing or cars, 21.6% are used for "bill consolidation" that is simply a reshuffling of liabilities, and 9.6% are used for education expenses. If durables are the most important consumption outlet for loans, then the consumption flow from durable purchases is approximately matched by the repayment schedule for the loan. If participants have only minimal liquid assets outside of their 401(k) (why would they need a loan otherwise?) and spend the entire amount of the loan (it is costly to withdraw more than one plans on spending because of the foregone tax benefits), then these repayments must be coming from either reducing the 401(k) contribution rate and hence increasing measured consumption, or from reducing other consumption flows by the amount of the loan repayments, leaving total consumption unchanged. Although the coordination between loan repayments and consumption from the loan is unlikely to be perfect—for example, the purchased durable could completely depreciate in three years, but the loan is paid off over five years—we believe that the measurement error from disregarding loan activity is second order. Moreover, the error induced by timing mismatches distorts the magnitude of MPC estimates, but it cannot explain why all of our MPC estimates are negative. To see this, suppose that a positive wealth shock induces an employee to take out a loan and purchase a durable that yields a constant consumption flow for three years and then ceases to exist. The loan is paid off over five years, and this repayment is funded by a decrease in the 401(k) contribution rate. Then in the year after the inception of the loan, the increase in consumption is underestimated by the change in the contribution rate. However, we would not measure an actual decrease in consumption through the contribution rate. Similar logic applies if the durable lasts for longer than the term of the loan; the consumption change is overestimated but does not take on the wrong sign.

E. Outside wealth shocks

The measures of wealth shock presented above are calculated only within the 401(k), but the theory calls for a comprehensive wealth shock measure. Most participants do not have significant financial wealth outside the 401(k), so this discrepancy is not an issue. For those who do have significant outside financial assets, to the extent that outside assets' return shocks are uncorrelated or positively correlated with the shocks within the 401(k), our MPC estimates will be unbiased or positively biased. Using typical returns to measure wealth shocks, we would expect a positive correlation if, within each asset class, individuals' investments outside the 401(k) are similar to those inside the 401(k). Using adaptive returns, we would expect a positive correlation if both inside and outside assets load positively on beta risk.

Even though 401(k) participants generally have few financial assets outside of their 401(k), many of them have a significant non-financial asset—owner-occupied housing—whose return we do not observe. In order for housing wealth shocks to materially affect our results, the surprise component of any real wealth innovation that occurs through housing must be negatively correlated with the surprise component of 401(k) wealth innovations. Intuitively, this seems unlikely to be true.

Nevertheless, to mitigate the possible effect of housing price changes, we conduct two tests. In Table 11, we present regression results where we have added dummy variables for the employee's state of residence interacted with year dummies to the baseline specification in equation (4). This allows us to control in a crude way for state-level variation in real estate price appreciation. In Table 12, we present regression results using the baseline specification of equation (4) but including only the one-third of our sample that live in zip codes with the lowest proportion of housing units that are owner-occupied, as measured by the 2000 U.S. Census.[12]

In both tables, all of our MPC point estimates are negative, five of the eight significantly so. Five of the eight lagged MPC estimates are also negative. Admittedly, our ability to control for housing wealth is limited because of the nature of our data. However, in those tests that we are able to conduct, we find no evidence that would call our central result into doubt.

F. Consumption adjustment is occurring through increased leisure

In response to a wealth shock, employees might adjust their consumption of leisure in ways that don't change expenditures and hence the 401(k) contribution rate. For example, an employee who has experienced a large positive wealth shock may decide to invest less effort at his job. This kind of effect is difficult to measure, but we may be able to infer its importance by examining another decision that is both observable and clearly increases leisure: the decision to retire.

We take employees who were at least 60 years old at year-end 1999, and we run a probit on the probability that these employees leave their company during calendar year 2000. The dependent variables are normalized wealth shock in 1999 and 1998, and age, log of salary, and log of tenure at year-end 1999. The results in Table 13 show no statistically significant effect of wealth shock on retirement.[13]

[12] The average owner-occupied housing fraction (weighted by employees living in the zip code) in the lower third of the sample is 60.4%, versus 74.8% for the entire sample.
[13] Our finding that wealth shocks do not affect retirement timing is not inconsistent with the literature, which has come to mixed conclusions on the question. See Hurd and Boskin (1984), Burtless (1986), Krueger and Pischke (1992), Holtz-Eakin, Joulfaian and Rosen (1993), and Imbens et al. (2001).

V. Conclusion

We have presented evidence that the short-run MPC out of orthogonal 401(k) capital gains shocks is negative, in violation of standard theory. Moreover, the magnitudes of some of these negative estimates are quite large. Because most of our wealth shocks come from equity return innovations, our results suggest a new explanation for the low covariance between equity returns and consumption: many investors are influenced by a reinforcement learning heuristic that leads high equity returns to encourage saving and discourage consumption, even when those capital gains are not useful for predicting future returns. This reinforcement effect attenuates and can even reverse the classical (income-effect) prediction of high-frequency positive covariance between equity returns and consumption growth. Thus, equity returns appear relatively riskless and the equity premium seems too high.

References

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15, pp. 47-68.

Benartzi, Shlomo, 2001. "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock." *Journal of Finance* 56, pp. 1747-1764.

Bodkin, Ronald, 1959. "Windfall Income and Consumption." *American Economic Review* 49, pp. 602-614.

Burtless, Gary, 1986. "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement." *Review of Economic Studies* 53, pp. 781-805.

Charness, Gary, and Dan Levin, 2003. "Bayesian Updating vs. Reinforcement and Affect: A Laboratory Study." Mimeo, University of California, Santa Barbara.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2003. "Employees' Investment Decisions About Company Stock." Forthcoming in Olivia S. Mitchell and Stephen P. Utkus, editors, *Pension Design and Decision-Making Under Uncertainty*.

Erev, Ido, and Alvin E. Roth. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88, pp. 848-881.

Holden, Sarah and Jack VanDerhei, 2001. "401(k) plan asset allocation, account balances, and loan activity in 2000." *Investment Company Institute Perspective* 7(5), pp. 1-27.

Holden, Sarah and Jack VanDerhei, 2003. "401(k) plan asset allocation, account balances, and loan activity in 2001." *Investment Company Institute Perspective* 9(2), pp. 1-15.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, 1993. "The Carnegie Conjecture: Some Empirical Evidence." *Quarterly Journal of Economics* 108, pp. 413-435.

Hurd, Michael D. and Michael J. Boskin, 1984. "The Effect of Social Security on Retirement in the Early 1970s." *Quarterly Journal of Economics* 99, 767-790.

Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote, 2001. "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players." *American Economic Review* 91, pp. 778-794.

Kreinin, Mordecai E., 1961. "Windfall Income and Consumption: Additional Evidence." *American Economic Review* 51, pp. 388-390.

Krueger, Alan B. and Jorn-Steffen Pischke, 1992. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics* 10, pp. 412-437.

Landsberger, Michael, 1966. "Windfall Income and Consumption: Comment." *American Economic Review* 56, pp. 534-540.

Ludvigson and Steindel. "How Important Is the Stock Market Effect on Consumption?" *Federal Reserve Bank of New York Economic Policy Review* 5, 29-52, July 1999.

Profit Sharing/401(k) Council of America, 2001. *44th Annual Survey of Profit Sharing and 401(k) Plans*. Chicago, IL: The Council.

Sevak, Purvi, 2002. "Weath Shocks and Retirement Timing: Evidence from the Nineties." University of Michigan Retirement Research Center Working Paper 2002-027.

Starr-McCluer, Martha, 2002. "Stock Market Wealth and Consumer Spending." *Economic Inquiry* 40, pp. 69-79.

Sundén, Annika, and Brian Surette, 2000. "Borrowing from 401(k) Plans: Implications for Retirement Wealth." *Just the Facts* 1.

U.S. Department of Labor, Bureau of Labor Statistics, 2003. *National Compensation Survey: Employee Benefits in Private Industry in the United States, 2000.* Washington, DC: U.S. Government Printing Office.

Table 1. Comparison of Companies in Study

Characteristic	Company A	Company B	Company C	Company D	Company E
Industry	Manufacturing	Healthcare	Manufacturing	Utility	Electronics
Number of employees	Over 20,000	Over 50,000	Over 20,000	Over 10,000	Over 10,000
Average age	44.1	42.7	44.6	43.5	39.5
Average salary	$51,835	$33,156	$66,700	$70,069	$54,702
% male	80%	19%	*	83%	65%
% married	56%	55%	75%	*	50%
Participation rate in 401(k) plan	80%	61%	86%	85%	83%
Average balance	$80,740	$19,501	$81,122	$88,033	$60,426
Maximum contribution rate (% of salary)	24%	15%	20%	25%	14%
Employer match features	Match first 6% of contribution. Match rate varies from 25% up to 100% by location.	Match 25% of first 3% of contribution	None until 2000. Afterwards, match 100% of first 1% of contribution, 50% of next 4% of contribution	Match 50% of first 7 or 8% of contribution, depending on union membership	Match 100% of first 3% of contribution, match 50% of next 3% of contribution
Investment funds	3 bond, 3 large-cap, 1 mid-cap, 1 small-cap, 3 overseas, company stock	1 cash, 1 bond, 3 pre-mix, 2 large-cap, 1 small-cap, 1 overseas, company stock	1 cash, 3 bond, 4 pre-mix, 7 large-cap, 5 mid-cap, 4 small-cap, 8 overseas, 3 sector, company stock	3 cash, 1 bond, 3 pre-mix, 1 large-cap, 1 mid-cap, 1 small-cap, 1 overseas, company stock	1 bond, 3 pre-mix, 5 large-cap, 2 overseas
Number of outstanding loans allowed	1 home loan, 1 general purpose loan	1	2	2	2
Non-hardship withdrawal rules before age 59½	1 withdrawal allowed per month from after-tax, rollover, company match, and profit-share balances	After-tax and employer contribution money from grandfathered plans can be withdrawn at any time	Not allowed	After-tax and employer match money can be withdrawn at any time	After-tax and rollover balances can be withdrawn at any time

* Data unavailable

Table 2. Indices Used to Calculate Adaptive Returns

This table presents the indices used in calculating the adaptive (ten-year lagged asset class) returns. All index returns assume that distributions are reinvested. The MSCI indices were obtained from the Morgan Stanley Capital International website. The Wilshire indices were obtained from the Wilshire Associates website. The 1-month T-bill returns were obtained from Kenneth French's website. The bond index was obtained from Datastream. S&P 500 and company stock data were obtained from finance.yahoo.com.

Asset Class	Index
Money Market	1-month T-bill return
GIC/Stable Value	1 month T-bill return
Bond	Lehman Brothers U.S. Aggregate Bond Index
Balanced	(mixed according to particular plan's funds)
Lifestyle/Pre-mix	(mixed according to particular plan's funds)
Large US Equity	S&P 500
Mid US Equity	Wilshire MidCap 500
Small US Equity	Wilshire SmallCap 1750
International	MSCI AC World Index Free
Emerging Markets	MSCI Emerging Markets Free*
Company Stock	Company Stock

* Inception date is December 1987. For dates when 10 years of data were not available, all the data available were used to calculate lagged return.

Table 3. Consumption Growth and Normalized Wealth Shock Distributions

This table presents summary statistics on consumption growth and two measures of wealth shock for the sample used in the tobit regressions in Table 4. Consumption growth is defined as the year-over-year change in consumption in the last two weeks of December, normalized by income in the last two weeks of the prior year. The normalized typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year and normalized by prior-year annual income. The normalized adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes, normalized by prior-year annual income. The distributions shown are for yearly consumption growth from year-end 1998 through year-end 2000, and wealth shocks from 1997 to 2000.

	Consumption growth	Normalized typical return wealth shock	Normalized adaptive wealth shock
Max	0.4251	5.2560	7.4919
99th percentile	0.2454	0.5152	1.3152
90th percentile	0.1292	0.0917	0.2289
75th percentile	0.0631	0.0178	0.0443
50th percentile	0.0102	-0.0022	-0.0003
25th percentile	-0.0161	-0.0322	-0.0858
10th percentile	-0.0567	-0.1223	-0.3063
1st percentile	-0.1548	-0.7533	-1.2075
Minimum	-0.2864	-4.7852	-8.7422
Mean	0.0249	-0.0159	-0.0135
Std. Dev.	0.0770	0.2050	0.4041

Table 4. Regression of Consumption Growth on Normalized Wealth Shock

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0049**	-0.0054**	-0.0045	-0.0165**
	(0.0015)	(0.0010)	(0.0024)	(0.0017)
$Shock_{t-1}/Y_{t-1}$	-0.0041**	0.0104**	-0.0012	-0.0169**
	(0.0011)	(0.0014)	(0.0021)	(0.0031)
Age_{t-1}	-0.0016**	-0.0015**		
	(0.0003)	(0.0003)		
Age^2_{t-1}	0.0000**	0.0000*	-0.0001**	-0.0002**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log(Y_{t-1})	0.0080**	0.0071**	-0.2097**	-0.2104**
	(0.0010)	(0.0010)	(0.0041)	(0.0041)
Log($Tenure_{t-1}$)	-0.0102**	-0.0108**	0.0391**	0.0406**
	(0.0005)	(0.0005)	(0.0022)	(0.0022)
σ	0.0770**	0.0770**		
	(0.0002)	(0.0002)		
N	69,676	69,676	54,560	54,560

* Significant at 5% level
** Significant at 1% level

Table 5. Regression of Consumption Growth on Normalized Wealth Shock Interacted with Age Dummy Variables

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. In the tobit regressions, the dummy variables ($n_1 \leq Age < n_2$) are equal to 1 if age at the end of $t - 26$ falls in the specified range; in the individual fixed-effects regressions, the variables are equal to 1 if age at year-end 1998 falls in the specified range. Y_{t-1} is inflation-adjusted salary paid in year $t - 1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for age, age-squared, log of tenure, log of salary, and company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t \times (Age < 30) / Y_{t-1}$	-0.0290	-0.0253*	-0.1025**	-0.0883**
	(0.0217)	(0.0108)	(0.0303)	(0.0171)
$Shock_t \times (30 \leq Age < 40) / Y_{t-1}$	-0.0094*	-0.0131**	-0.0197**	-0.0349**
	(0.0045)	(0.0027)	(0.0074)	(0.0043)
$Shock_t \times (40 \leq Age < 50) / Y_{t-1}$	-0.0068**	-0.0067**	-0.0059	-0.0168**
	(0.0023)	(0.0014)	(0.0035)	(0.0023)
$Shock_t \times (50 \leq Age < 60) / Y_{t-1}$	-0.0017	-0.0033*	0.0015	-0.0094**
	(0.0023)	(0.0013)	(0.0035)	(0.0027)
$Shock_t \times (Age \geq 60) / Y_{t-1}$	-0.0017	-0.0054	-0.0079	-0.0055
	(0.0085)	(0.0052)	(0.0148)	(0.0111)
$Shock_{t-1} \times (Age < 30) / Y_{t-1}$	-0.0254	0.0509*	-0.0538	-0.1108**
	(0.0157)	(0.0201)	(0.0301)	(0.0346)
$Shock_{t-1} \times (30 \leq Age < 40) / Y_{t-1}$	0.0020	0.0195**	-0.0134	-0.0331**
	(0.0033)	(0.0045)	(0.0078)	(0.0086)
$Shock_{t-1} \times (40 \leq Age < 50) / Y_{t-1}$	-0.0024	0.0116**	-0.0022	-0.0151**
	(0.0016)	(0.0021)	(0.0032)	(0.0043)
$Shock_{t-1} \times (50 \leq Age < 60) / Y_{t-1}$	-0.0055**	0.0092**	0.0020	-0.0098*
	(0.0015)	(0.0019)	(0.0029)	(0.0048)
$Shock_{t-1} \times (Age \geq 60) / Y_{t-1}$	-0.0076*	0.0111	-0.0054	-0.0050
	(0.0030)	(0.0063)	(0.0067)	(0.0175)
σ	0.0770**	0.0770**		
	(0.0002)	(0.0002)		
N	69,676	69,676	54,560	54,560

* Significant at 5% level

** Significant at 1% level

Table 6. Regression of Consumption Growth on Normalized Wealth Shock for Participants Interior to Match Threshold

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. In the tobit regressions, we exclude those individuals whose $t-1$ contribution rate is greater than or equal to the threshold to which the employer will match contributions. In the employee fixed-effects regressions, we exclude those individuals whose year-end 1998 contribution rate is greater than or equal to the threshold to which the employer will match contributions. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0009	-0.0004	-0.0245**	-0.0143**
	(0.0066)	(0.0040)	(0.0092)	(0.0037)
$Shock_{t-1}/Y_{t-1}$	-0.0035	0.0176**	-0.0082	0.0012
	(0.0050)	(0.0057)	(0.0082)	(0.0067)
Age_{t-1}	-0.0022**	-0.0022**		
	(0.0008)	(0.0008)		
Age^2_{t-1}	0.0000*	0.0000*	-0.0001	-0.0001
	(0.0000)	(0.0000)	(0.0001)	0.0000
Log(Y_{t-1})	-0.0075**	-0.0085**	-0.3204**	-0.3203**
	(0.0028)	(0.0028)	(0.0111)	(0.0079)
Log($Tenure_{t-1}$)	-0.0136**	-0.0140**	0.0329**	0.0341**
	(0.0014)	(0.0014)	(0.0043)	(0.0030)
σ	0.0787**	0.0786**		
	(0.0006)	(0.0006)		
N	11,113	11,113	8,896	8,896

* Significant at 5% level

** Significant at 1% level

Table 7. Regression of Consumption Growth on Normalized Wealth Shock for Participants Who Remain in the Same Zip Code

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. In the tobit regressions, we exclude those individuals whose residential zip code changes between $t-1$ and t. In the employee fixed-effects regressions, we exclude those individuals whose residential zip code changes any time between year-end 1998 and year-end 2000. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0045**	-0.0048**	-0.0057*	-0.0164**
	(0.0015)	(0.0010)	(0.0024)	(0.0018)
$Shock_{t-1}/Y_{t-1}$	-0.0043**	0.0104**	-0.0019	-0.0171**
	(0.0011)	(0.0014)	(0.0021)	(0.0031)
Age_{t-1}	-0.0014**	-0.0013**		
	(0.0003)	(0.0003)		
Age^2_{t-1}	0.0000*	0.0000	-0.0001**	-0.0001**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log(Y_{t-1})	0.0067**	0.0057**	-0.2047**	-0.2055**
	(0.0010)	(0.0011)	(0.0043)	(0.0043)
Log($Tenure_{t-1}$)	-0.0103**	-0.0109**	0.0387**	0.0403**
	(0.0005)	(0.0005)	(0.0024)	(0.0024)
σ	0.0757**	0.0757**		
	(0.0002)	(0.0002)		
N	64,735	64,735	47,870	47,870

* Significant at 5% level

** Significant at 1% level

Table 8. Regression of Consumption Growth on Normalized Wealth Shock for Unmarried Participants

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. In the tobit regressions, we exclude those individuals who were married as of the end of $t-1$. In the employee fixed-effects regressions, we exclude those individuals who were married at any time between year-end 1998 and year-end 2000. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0078*	-0.0111**	-0.0098*	-0.0220**
	(0.0035)	(0.0023)	(0.0047)	(0.0034)
$Shock_{t-1}/Y_{t-1}$	-0.0042*	0.0168**	-0.0023	-0.0202**
	(0.0019)	(0.0034)	(0.0035)	(0.0063)
Age_{t-1}	-0.0012*	-0.0011*		
	(0.0005)	(0.0005)		
Age^2_{t-1}	0.0000	0.0000	-0.0002*	-0.0002**
	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Log(Y_{t-1})	-0.0029	-0.0042*	-0.2376**	-0.2378**
	(0.0020)	(0.0020)	(0.0075)	(0.0075)
Log($Tenure_{t-1}$)	-0.0101**	-0.0109**	0.0415**	0.0431**
	(0.0009)	(0.0009)	(0.0036)	(0.0036)
σ	0.0827**	0.0826**		
	(0.0004)	(0.0004)		
N	24,256	24,256	17,784	17,784

* Significant at 5% level

** Significant at 1% level

Table 9. Regression of Yearly Consumption Change, Assuming Immediate Consumption of In-Service Withdrawals, on Normalized Wealth Shock

The dependent variable is year-over-year change in annual consumption, normalized by annual salary in the prior year. We assume that withdrawals from the 401(k) net of rollovers are consumed immediately. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0012	0.0002	0.0010	-0.0182**
	(0.0016)	(0.0010)	(0.0027)	(0.0020)
$Shock_{t-1}/Y_{t-1}$	-0.0054**	0.0068**	-0.0024	-0.0181**
	(0.0011)	(0.0015)	(0.0024)	(0.0035)
Age_{t-1}	-0.0010**	-0.0009**		
	(0.0003)	(0.0003)		
Age^2_{t-1}	0.0000	0.0000	-0.0003**	-0.0003**
	0.0000	0.0000	0.0000	0.0000
Log(Y_{t-1})	-0.0047**	-0.0053**	-0.9525**	-0.9518**
	(0.0010)	(0.0010)	(0.0097)	(0.0097)
Log($Tenure_{t-1}$)	-0.0082**	-0.0086**	0.0456**	0.0594**
	(0.0005)	(0.0005)	(0.0071)	(0.0072)
σ	0.0786**	0.0786**		
	(0.0002)	(0.0002)		
N	68,283	68,283	53,439	53,439

* Significant at 5% level

** Significant at 1% level

**Table 10. Regression of Yearly Consumption Change,
Assuming Annuitization of In-Service Withdrawals, on Normalized Wealth Shock**

The dependent variable is year-over-year change in annual consumption, normalized by annual salary in the prior year. We assume that 5% of a withdrawal from the 401(k) net of rollovers is consumed each year starting in the year of the withdrawal. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0020	0.0000	0.0012	-0.0152**
	(0.0014)	(0.0009)	(0.0018)	(0.0014)
$Shock_{t-1}/Y_{t-1}$	-0.0044**	0.0072**	0.0004	-0.0118**
	(0.0010)	(0.0013)	(0.0016)	(0.0024)
Age_{t-1}	-0.0008**	-0.0007**		
	(0.0003)	(0.0003)		
Age^2_{t-1}	0.0000	0.0000	-0.0003**	-0.0003**
	0.0000	0.0000	0.0000	0.0000
Log(Y_{t-1})	-0.0022*	-0.0029**	-0.9747**	-0.9747**
	(0.0009)	(0.0010)	(0.0067)	(0.0067)
Log($Tenure_{t-1}$)	-0.0081**	-0.0085**	0.0558**	0.0696**
	(0.0005)	(0.0005)	(0.0049)	(0.0050)
σ	0.0727**	0.0727**		
	(0.0002)	(0.0002)		
N	69,676	69,676	54,560	54,560

* Significant at 5% level
** Significant at 1% level

Table 11. Regression of Consumption Change on Normalized Wealth Shock with State × Year Effects

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. The regressions include year dummies interacted with dummies for state of residence, the coefficients of which we do not report. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0047**	-0.0057**	-0.0035	-0.0142**
	(0.0015)	(0.0010)	(0.0024)	(0.0017)
$Shock_{t-1}/Y_{t-1}$	-0.0044**	0.0099**	-0.0010	-0.0135**
	(0.0011)	(0.0014)	(0.0021)	(0.0031)
Age_{t-1}	-0.0015**	-0.0014**		
	(0.0003)	(0.0003)		
Age^2_{t-1}	0.0000*	0.0000*	-0.0001**	-0.0002**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log(Y_{t-1})	0.0057**	0.0049**	-0.2013**	-0.2022**
	(0.0010)	(0.0010)	(0.0041)	(0.0041)
Log($Tenure_{t-1}$)	-0.0081**	-0.0086**	0.0389**	0.0403**
	(0.0005)	(0.0005)	(0.0022)	(0.0022)
σ	0.0760**	0.0759**		
	(0.0002)	(0.0002)		
N	69,676	69,676	54,560	54,560

* Significant at 5% level

** Significant at 1% level

Table 12. Regression of Consumption Growth on Normalized Wealth Shock for Participants in the Lowest Third of Home Ownership Probability Distribution

The dependent variable is year-over-year change in consumption in the last two weeks of December, normalized by salary in the last two weeks of the prior year. The first two columns present coefficients from a tobit estimation, and the second two columns present coefficients from a least-squares regression with employee fixed effects. In the tobit regressions, we exclude those individuals who, at the end of year $t-1$, lived in a zip code whose owner-occupied housing units as a fraction of total housing units is in the top two-thirds of our sample. In the employee fixed-effects regressions, we exclude those individuals who, at the end of 1998, lived in a zip code whose owner-occupied housing units as a fraction of total housing units is in the top two-thirds of our sample. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. σ is the estimated standard deviation of the latent variable's disturbance term in the tobit regressions. Coefficients for company-year dummies are omitted. Values in parentheses are standard errors.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0057	-0.0064**	-0.0054	-0.0165**
	(0.0032)	(0.0020)	(0.0044)	(0.0031)
$Shock_{t-1}/Y_{t-1}$	-0.0063**	0.0164**	0.0006	-0.0146**
	(0.0020)	(0.0029)	(0.0034)	(0.0054)
Age_{t-1}	-0.0024**	-0.0024**		
	(0.0006)	(0.0006)		
Age^2_{t-1}	0.0000**	0.0000**	-0.0001	-0.0001
	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Log(Y_{t-1})	-0.0002	-0.0018	-0.2566**	-0.2568**
	(0.0020)	(0.0020)	(0.0076)	(0.0076)
Log($Tenure_{t-1}$)	-0.0091**	-0.0099**	0.0430**	0.0446**
	(0.0010)	(0.0010)	(0.0038)	(0.0038)
σ	0.0872**	0.0872**		
	(0.0005)	(0.0005)		
N	23,696	23,696	17,494	17,494

* Significant at 5% level

** Significant at 1% level

Table 13. Probit Regression of Older Workers Leaving the Company on Normalized Wealth Shock

The dependent variable equals 1 if the participant has left the company by year-end 2000, and 0 otherwise. Individuals included in these regressions were active participants in their company's 401(k) plan for all of calendar years 1998 and 1999 and were at least 60 years old at year-end 1999. $Shock_t$ is defined in two ways. The typical return wealth shock is the difference between the capital gain actually realized by the participant in a fund and what he would have realized had he held the same amount in a representative basket of mutual funds in the fund's asset class, summed over all funds and months in a year. The adaptive wealth shock is the difference between the capital gain actually realized by the participant and what he would have realized had the returns on his funds equaled the ten-year lagged average returns of the funds' respective asset classes. Y_{t-1} is inflation-adjusted salary paid in year $t-1$. Age_{t-1} is the participant's age at the end of $t-1$. $Tenure_{t-1}$ is the number of years since original hire at the end of $t-1$. Coefficients for company dummies are excluded. Values in parentheses are standard errors.

	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	0.1911	0.0537
	(0.1380)	(0.0949)
$Shock_{t-1}/Y_{t-1}$	-0.0793	-0.0533
	(0.1544)	(0.4337)
Age_{t-1}	0.0377*	0.0367
	(0.0191)	(0.0191)
$\text{Log}(Y_{t-1})$	-0.6297**	-0.6609**
	(0.1619)	(0.1648)
$\text{Log}(Tenure_{t-1})$	0.3850**	0.3722**
	(0.0951)	(0.0954)
N	1,419	1,419

* Significant at 5% level
** Significant at 1% level