

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: April 12, 2004

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, Matthew Shapiro, Bill Zame, and from the comments of seminar participants at Berkeley, Harvard, the NBER Monetary Economics Program Meeting, UCLA and RAND. 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, National Institute on Aging Grant T32-AG00186, and the Mustard Seed Foundation. 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 shortly after the windfall is received. We test this prediction 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 *idiosyncratic* 401(k) capital gains shocks. These results cannot be interpreted as standard intertemporal substitution, since the idiosyncratic returns that we study do *not* predict future returns. Instead, our findings imply that many investors are influenced by a positive feedback effect, through which higher recent returns encourage higher short-run saving. Like any other animal, 401(k) participants appear to increase behaviors that have been associated with high rewards in the past.

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, Suite 1400
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

The income effect predicts that an unexpected wealth windfall should immediately and permanently increase consumption. As an illustrative example, the marginal propensity to consume (MPC) out of windfalls roughly equals the discount rate in an infinite-horizon model with complete markets.¹ If markets are not complete, the MPC may considerably exceed the discount rate (Carroll 1992, 1997).

In this paper, we present empirical evidence that the short-run MPC with respect to 401(k) capital gain windfalls may actually be *negative*. At year-end 2001, 401(k) accounts contained \$1.75 trillion owned by 45 million workers (Holden and Vanderhei 2003). Using administrative data on 40,000 401(k) accounts, we calculate idiosyncratic capital gains in the cross-section that do *not* forecast future returns. Standard theory predicts that a cross-sectional regression of consumption growth on these orthogonal capital gains shocks will identify only a positive income effect, since the standard substitution effect should be inoperative. Contrary to theory, we estimate a negative short-run MPC.

At first glance, this result appears to be quite perverse. However, our results represent a natural generalization of an old idea: investors chase returns. Several authors have shown that return-chasing influences asset allocation (e.g., Patel, Zeckhauser, and Hendricks 1994, Benartzi 2001, Goetzmann and Massa 2004). We are the first to show that return-chasing also plays a role in *consumption* decisions.² High transitory past investment returns encourage (naïve) households to invest more in the short-run and therefore consume less.

Although the main thrust of our paper is empirical, we also show that combining reinforcement learning with the standard economic model can reconcile our results with the positive MPC estimates found in prior consumption studies. Reinforcement learning models have had success in predicting subject choices in experiments (Roth and Erev 1995, Erev and Roth 1998, Charness and Levin 2003). In our context, reinforcement learning encourages saving in response to high returns and discourages saving in response to low returns. Even though the

¹ In a continuous-time model with isoelastic preferences the MPC is given by $(1-\theta)[r + (\theta\pi^2)/(2\sigma^2)] + \theta\rho$, where θ is the elasticity of intertemporal substitution, r is the risk free real interest rate, π is the equity premium, σ is the standard deviation of stock returns, and ρ is the discount rate. When the elasticity of intertemporal substitution is one (log utility, a standard calibration case), this formula collapses to the discount rate.

² The first suggestions of such an effect can be found in work by Starr-McCluer (2002), who analyzes household surveys in which 11.6% of stockholders report that the 1990s bull market caused them to save more, while only 3.4% say they saved less. Our identification strategy differs from that of Starr-McCluer, since we identify wealth shocks cross-sectionally.

standard income effect is operative, it is more than offset by the reinforcement effect in this particular case, yielding a negative MPC.

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 the previous literature. 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) that 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 years and average salary of \$45,656 among full-time workers in companies employing over 1,000 workers and offering 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. All of the plans allow participants to take hardship withdrawals from and loans against their 401(k) plan balances, and two allow non-hardship

withdrawals from current contributions. These provisions make 401(k) savings in the companies we study more liquid than for the 401(k) participant population at large.³ Finally, all of the plans allow changes to the elected contribution rate and asset allocation on a daily basis. Changes can be made by talking to a benefits center representative on the phone during business hours, or by using a touch-tone phone system or the Internet 24 hours a day. The direct transaction costs involved in changing one's savings rate in these plans are therefore minimal.

II. Empirical methodology

Our empirical objective is to estimate the relationship between an orthogonal wealth shock in year t and consumption growth from year-end $t - 1$ to year-end t . The key assumption underlying our analysis is that consumption adjustments are observable through changes in the 401(k) contribution rate. This assumption is plausible: most households should be doing nearly all of their saving in the 401(k) because the employer match (for the plans in our data) and tax benefits make the 401(k) plan the most attractive savings vehicle available. Consistent with this normative view, among 401(k)-holding households earning between \$20,000 and \$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).⁴ 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 since the generosity of our 401(k) plans' early withdrawal and loan provisions substantially mitigates the need to maintain a precautionary stock of wealth outside the 401(k). We will, however, 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

³ 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.

⁴ 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.

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

where C_t is annualized consumption flow during the last pay cycle⁵ 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 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

$$C_t = Y_t(1 - S_t) + k, \quad (2)$$

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

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

We add the before-tax and (if the plan offers the option) after-tax 401(k) contribution rates in effect for the last pay cycle of 1998, 1999, or 2000 to calculate S_t in each of these years. Note that there are at most three consumption observations per individual, which translates into a maximum of two consumption growth observations for each employee. Y_t is set to annualized salary in the last pay cycle of year t in January 1998 dollars, assuming that salary during a calendar year is paid in nominally equal amounts each pay cycle and deflating by the monthly CPI series.

Constructing our key independent variable, the 401(k) wealth shock, is more complicated. Simply calculating the annual change in 401(k) balances would not generate a valid measure of the type of wealth shock that we are trying to identify. Much of the annual change in account balances is generated by payroll contributions to the plan, and these contributions should not be counted as wealth shocks because they are merely transfers from human wealth to

⁵ Our year-end data include the 401(k) contribution rate that applies to the final pay period of the year, which in our sample is typically two weeks.

financial wealth.⁶ We could subtract annual 401(k) contributions from the change in 401(k) balances to isolate the change in balances that is due to capital gains (or losses). But a portion of these capital gains is expected and thus, in theory, should not affect consumption growth. We must therefore isolate the *unexpected* component of capital gains. With these considerations in mind, we define the wealth shock accrued over a year as the difference between the realized and expected dollar denominated capital gain in each asset, summed over all assets.⁷

We calculate capital gains 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, and then add up across months and assets.⁸ In order to compute expected capital gains, we must measure expected percent returns. We use two different proxies. The first measure, the “typical return,” is defined for two asset classes: equities (including employer stock) and bonds. 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.^{9,10} 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. The typical return is motivated by the intuition that the deviation of a mutual fund's return from the contemporaneous average return in its asset class—the “typical” return—is unforecastable.

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.¹¹ Graham and Harvey (2003) and Vissing-Jorgensen (2003) report survey evidence that stock

⁶ More explicitly, human wealth can be seen as the present value of all future labor income. Each paycheck is a dividend that diminishes human wealth by the same amount it increases financial wealth. The employee can choose to consume this dividend and diminish total wealth, or save it and keep total wealth unchanged.

⁷ We have also estimated the regressions simply using realized capital gains as the wealth shock measure, and our results are qualitatively unchanged.

⁸ 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 .

⁹ The average is taken over participants with holdings in the relevant asset class. Returns on lifestyle/pre-mixed or balanced funds, which hold pre-specified proportions of equities and bonds, are not used to calculate the typical returns. However, typical equity and bond returns are combined in the appropriate proportions to form expected returns for lifestyle/pre-mixed and balanced funds.

¹⁰ If the plan offers only one fund in a particular asset class—not an unusual situation for bond funds—then expected returns and realized returns will always be the same in that asset class and the wealth shock arising from holding this asset class is always be zero. It is for this reason that we do not compute typical returns for cash funds; no plan in our data offers more than one cash fund.

market return expectations comove positively with recent market returns. Benartzi (2001) presents survey evidence that expectations of employer stock returns are highly correlated with lagged employer stock returns. Vissing-Jorgensen (2003), Benartzi (2001), and Choi et al. (2003) show that asset allocation patterns are consistent with these stated beliefs. 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.

To illustrate how our wealth shocks are calculated, suppose that there are only two funds available to investors. Fund A's expected return is 1% this month, and Fund B's expected return is 2%. Fund A actually returns 3%, and Fund B actually returns -3%. Investor 1 had \$3,000 in Fund A and \$1,000 in Fund B at the end of the prior month; Investor 2 had \$1,000 in Fund A and \$1,000 in Fund B. Then Investor 1 receives a $\$3,000 \times (3 - 1)\% + \$1,000 \times (-3 - 2)\% = \10 wealth shock this month, while Investor 2 receives a $\$1,000 \times (3 - 1)\% + \$1,000 \times (-3 - 2)\% = -\30 wealth shock this month. We sum monthly wealth shocks to form the annual wealth shock for each investor.

The identifying variation in our wealth shocks comes from both cross-sectional variation in unexpected fund returns and cross-sectional variation in the dollar amounts individuals invested in each fund at the beginning of the period. It is important to recognize that our estimation relies upon the size of an individual's wealth shock *relative* to other individuals' contemporaneous wealth shocks. This means that the identifying variation in the wealth shock is completely uninformative about future returns; in the previous paragraph's example, Investor 1 should make no inference about future returns from the fact that his wealth shock is \$40 higher than Investor 2's wealth shock.

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 may shift predictably with demographic variables, causing changes in consumption expenditures. 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, funds idiosyncratically outperform or underperform in sample, so our wealth shock measures may be correlated with other variables

¹¹ We have also run our regressions using one-year lagged returns, and our results are qualitatively unchanged.

that predict both consumption growth and fund choices. In light of these considerations, we modify (3) to arrive at our final regression equation,

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

where \mathbf{X}_{t-1} is a vector of demographic variables intended to control for demographically-related consumption shifts. We discuss the variables included in \mathbf{X}_{t-1} at the end of this section.

We take two different approaches to estimating equation (4). First, we stack our data across time periods and use a tobit to account for the fact that 401(k) contribution rates cannot be negative or exceed 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.¹²

The sample for our tobit regressions is limited to employees who have been actively employed at the firm and enrolled in the 401(k) plan for at least two complete consecutive calendar years between 1998 and 2000. We include employees whose contribution rate or plan balances are zero, provided that they had positive balances at some time in the past. We also require that individuals have salaries greater than \$20,000 in year $t - 1$ because a large fraction of those with salaries under \$20,000 are part-time employees.¹³ In addition, we 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.¹⁴ 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 year $t - 1$ that they could exceed the \$10,000 annual limit on before-tax 401(k) contributions by contributing at the plan's

¹² A natural extension of our empirical strategy would be to estimate a fixed-effects tobit. Because of an incidental parameters problem, such an estimation is not possible using maximum likelihood methods. Honoré (1992) presents a semiparametric approach to estimating censored models with fixed effects. The fact that our data are censored at both the left and the right necessitates the use of the least absolute deviations version of Honoré's estimator. Unfortunately, there is considerable theoretical ambiguity about how standard errors for this estimator should be computed. Therefore, we have not pursued this avenue.

¹³ In the March 2001 Current Population Survey, 29.9% of workers who earned less than \$20,000 a year worked less than 35 hours a week or fewer than 40 weeks per year. Only 5.6% of workers earning between \$20,000 and \$30,000 a year satisfied this definition of part-time work.

¹⁴ 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.

maximum before-tax 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 before-tax dollar limit midway through the year. If his company only allows before-tax contributions (Companies B, C, and E), then his 401(k) contribution rate is frozen at 0 for the remainder of the year and does not reflect any changes in his consumption rate. If the company allows additional after-tax contributions to the 401(k) (Companies A and D) or before-tax contributions to a non-401(k) deferred compensation plan (Company A), consumption changes may still not be reflected in the contribution rate because these alternative savings vehicles are not as attractive and therefore may not attract the marginal savings dollar.¹⁵

These criteria leave us with a final tobit sample of 69,581 observations on 42,554 employees. The typical returns for the tobit regressions are estimated by averaging over this sample.

For our regressions incorporating employee-level fixed effects, 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 continuously enrolled in the 401(k) plan from January 1, 1998 to December 31, 2000. The cutoffs we use to trim salary growth outliers and workers with high salaries in 1998 remain the same as in the original tobit sample. This leaves us with 54,054 observations on 27,027 employees in the fixed-effects sample. We estimate a separate set of typical returns for the fixed-effects regressions using this sample.¹⁶

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 year-end $t - 1$, in the demographic vector \mathbf{X}_{t-1} . Additionally, we include company-year dummies. When running the employee fixed-effects regressions, we must change the set of explanatory variables. 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.

¹⁵ Deferred compensation plans are subject to bankruptcy risk; if the employer declares bankruptcy, the employee can lose his entire deferred compensation balance. We eliminate participants in Company A's deferred compensation plan from our sample.

¹⁶ We have run the tobit regressions on the fixed effects sample and found no substantive differences in our results.

III. Main results

A. Regression estimates

Table 3 reports summary statistics for our consumption growth and normalized wealth shock measures in the tobit sample. Consumption growth has a median of 1.0% of income, exhibits slight positive serial correlation, and 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 11.9% of a year's income, and the 90th percentile is a gain of 8.9% of a year's income. The spread is wider for the normalized adaptive wealth shock; the 10th and 90th percentiles are separated by 50.4% of annual income. We will show that this amount of variation in wealth shocks is sufficient to produce precise estimates of the MPC. Both normalized wealth shocks are *ex post* negatively serially correlated. This serial correlation is sample-specific and arises from the persistence of individual asset allocations and the fact that the stock market did well in the first half of our sample and poorly in the second half. This should not be interpreted as evidence that current wealth shocks are informative about future wealth shocks.¹⁷

Table 4 presents the coefficients from estimating equation (4). The first row shows that all four MPC estimates are negative and statistically significant. The point estimates range from -0.37% to -1.82% , which means that a positive wealth shock equal to one year of income will contemporaneously decrease annualized consumption flow by 0.37% to 1.82% of yearly income. There is no consistent evidence that consumption increases with a lag in response to positive wealth shocks. The second row shows that three of the four lagged response estimates are negative, and the only significantly positive estimate of 1.05% is counterbalanced by significantly negative estimates of -0.53% and -1.85% .

The reader may be concerned that these results are not driven by active savings responses to wealth innovations, but rather reflect completely passive participants whose income growth happens to be correlated with wealth shocks in sample.¹⁸ To address this concern, we have also estimated the main regressions setting the income adjustment term in the dependent variable to zero, so that we are only measuring active contribution rate changes, $S_{t-1} - S_t$. (These results are

¹⁷ If participants who held large amounts of equity in the first half of the sample and hence had high measured wealth shocks foresaw that equity returns would be low in the second half of the sample, they could have easily traded out of equities and thus avoided the negative serial correlation in wealth shock.

not reported in the tables.) All four of our contemporaneous MPC estimates remain negative, two are statistically significant at the 1% level, and one is statistically significant at the 5% level.

The certainty-equivalence lifecycle 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. This motivates our estimation of 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.

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 -3.04% for employees under 30 and an MPC of -0.38% for employees over 60. The corresponding employee fixed effects regression yields an estimate of -11.33% for those under 30 and -0.87% for those over 60. The lagged MPC does not exhibit a consistent 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 MPCs are restricted to a particular subset of the population.¹⁹ If our results simply reflected reallocation between outside assets and the 401(k), or if outside wealth shocks were contaminating our estimates, then we would expect to find positive MPC estimates among the young, who are the least likely to own non-401(k) assets. Instead, the young have the most negative MPCs. If negative MPCs were simply a pathology of inexperienced investors that disappears in the long run, we would expect to see positive MPC estimates among the old. But in our data, even the oldest employees have negative MPCs.

In the long run, consumption and wealth are tied together by the budget constraint. If in the short run, 401(k) investors' consumption growth comoves negatively with wealth shocks, when does consumption catch up to wealth? A plausible answer is that consumption moves

¹⁸ See Samuelson and Zeckhauser (1988), Madrian and Shea (2001), and Choi et al. (2002 and 2004) for evidence on participant inertia in 401(k) plans.

¹⁹ 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* effect on contemporaneous MPC, while age continues to have a positive effect. There is no consistent pattern for the interaction of age and salary with $Shock_{t-1}$.

towards wealth in retirement, either through a higher consumption rate during retirement or an earlier retirement date. Table 5 presents evidence that suggests that at least the latter is happening. We take employees who were at least 60 years old at year-end 1999 and run a probit on the probability that these employees left the company in 2000. (Overall, 50.4% of this population left in 2000.) The dependent variables are normalized wealth shock in 1999 and 1998, company dummies, and age, log of salary, and log of tenure at year-end 1999. Using typical returns, positive wealth shocks significantly increase the probability of leaving; at the mean of the explanatory variables, a one standard deviation increase in the 1999 wealth shock raised the probability of leaving by 2.7 percentage points. Using adaptive returns, there is no statistically significant wealth effect, but the point estimates are economically large; a one standard deviation increase in the 1999 wealth shock raised the probability of leaving by 1.1 percentage points, and a one standard deviation increase in the 1998 wealth shock raised the probability of leaving by 5.1 percentage points. These results are consistent with the findings of Hurd and Boskin (1984), Holtz-Eakin, Joulfaian and Rosen (1993), Imbens et al. (2001), and Sevak (2003). (See, however, Burtless (1984) and Krueger and Pischke (1992), who find no wealth effect on labor supply.)

B. Interpretation and Reconciliation with Past MPC Estimates

Our estimates are in sharp contrast to past empirical research that has found positive MPCs out of cash windfalls and aggregate stock market movements. We square our results with these previous findings by interpreting negative MPCs as the consequence of reinforcement learning working in opposition to the standard income effect. Roth and Erev (1995) identify two key features of reinforcement learning models. The first is the Law of Effect: agents are more likely to repeat actions that have yielded favorable outcomes in the past. The second is the Power Law of Practice: learning curves are steep initially and then level out as the stock of reinforcements increases.

Reinforcement learning reconciles all of the evidence on MPCs:

- The Law of Effect predicts that if a consumer experiences a gratifyingly high return from her savings activity, she will allocate more resources to savings and less to consumption. This behavior is a generalization of the phenomenon commonly referred to as returns-chasing, where investors reallocate money to assets that have recently experienced high returns while holding total savings fixed. The Law of

Effect is offset by the income effect, which pushes the investor to cut her savings. In our data, the reinforcement effect appears to dominate the income effect, causing the MPC to be negative even though the returns from which the MPC is identified convey no information about future returns.

- The Power Law of Practice is consistent with our finding that the young—whose stock of past reinforcements is the smallest—have the most negative MPCs, since they are the most swayed by the Law of Effect and have the smallest income effect.
- When financial investment does not causally precede windfalls, there is no reinforcement to invest more and only the income effect is operative. Hence, the MPC out of war veteran payments and Holocaust reparations range from 17% to 97% (Bodkin 1959, Kreinin 1961, Landsberger 1966).
- The income effect increases proportionally with the size of the windfall, whereas the reinforcement effect levels off due to the Power Law of Practice. Therefore, the MPC out of large lottery prizes (Imbens et al. 2001) is 86%.²⁰
- Reinforcement learning—and hence negative MPCs—should be less prevalent among agents who are more sophisticated than the 401(k) investors in our sample. We also expect that the sophisticated are, on average, richer than the unsophisticated. Therefore, aggregate consumption growth comoves positively with aggregate stock market movements because sophisticates do most of the consuming in the economy. Consistent with this hypothesis, Starr-McCluer (2002) finds that households with more than \$250,000 in stockholdings were much more likely than poorer stockholders to report that the 1990s bull market caused them to increase their spending.

IV. Robustness checks

In this section, we further 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. We find no evidence that weakens the force of our main result. Of the 28 MPC estimates presented in this section, 26 are negative, and 15 of these

²⁰ Kearney and Liao (2004) do not find evidence of a negative same-day MPC out of small lottery prizes. However, they do not measure the effect of prizes on lottery purchases on subsequent days.

are statistically significant. None of the two positive point estimates are statistically significant, and they are each an economically negligible 0.04%.

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 risk-free returns. Therefore, these employees have especially strong incentives to adjust their consumption expenditures exclusively through their 401(k) contribution rate.²¹ Because Company C did not have a match until 2000, its participants are excluded from this analysis.

The results of this regression are found in Table 7. The tobit estimates of the MPC are attenuated but still negative at -0.11% and -0.01% . In contrast, the two employee fixed effect estimates are more negative than in the baseline regression and statistically significant at -2.40% and -1.57% . Among the lagged MPC estimates, three out of the four are negative.

Participants may also respond to positive wealth shocks by purchasing a home. Because mortgage payments cover both the rental flow of housing services and the purchase of home equity, the home becomes a savings vehicle. The amount of money that goes towards home equity could 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

²¹ 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

Table 8 show that all four MPC estimates are negative, three are statistically significant, and the fourth barely misses significance ($p = 0.06$). 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. Alternatively, the consumption changes we measure may only reflect reshuffling of assets between the two 401(k) accounts. 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 9. We find that all four MPC estimates are negative: two of them are significant at the 1% level, one at the 5% level, and one at the 10% level. 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 4.9% 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 tax-deferred 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.

incentives to save in the 401(k). However, it turns out that only 8.0% of those in the restricted tobit sample are not vested at all at year-end 1998, while 81.2% are fully vested (the rest are partially vested).

We estimate 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.²² 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) 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 -18.5% and 34.7%. We do not trim the dependent variable in the regressions that assume that the withdrawals are annuitized.

Table 10 presents the results when we assume that net withdrawals are consumed immediately, and Table 11 presents the results when we assume that net withdrawals are

²² 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.

annuitized. 6 out of the 8 MPC estimates are negative, including both statistically significant estimates. The magnitude of the two positive estimates is an economically trivial 0.04%. 2 out of the 8 lagged MPC estimates are negative, and 2 out of the 5 significant lagged 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 a 401(k) loan for up to 50% of vested plan balances or \$50,000, whichever is lower. 401(k) loans are typically repaid through payroll deductions, with interest accruing to the individual's own 401(k) account. The maximum term of a 401(k) loan varies with its purpose. Legally, a primary residence loan can have a term of up to 30 years, while a general purpose loan can have a term of no longer than 5 years. In practice, many companies further restrict the term of the loans that they offer their 401(k) participants. Early repayment is possible with no penalties. While many companies charge loan origination fees, none of the companies in our data choose to do so. Loans are common in our sample: 29.0% of the tobit sample had a loan outstanding at some point between year-ends 1998 and 2000.

Although a significant fraction of our population has a 401(k) loan, 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, while the loan is paid off over five years—we believe that the measurement error from disregarding loan activity is of 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 the wrong sign.

E. Outside wealth shocks

The measures of wealth shock presented above are calculated only within the 401(k), but theory calls for a comprehensive wealth shock measure. Most participants—particularly the young, for whom we estimate the most negative MPCs—do not have significant financial wealth outside the 401(k). 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 qualitatively 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. This seems unlikely to be true, particularly when considering typical return wealth shocks.

Nevertheless, to mitigate the possible effect of housing price changes, we conduct two tests. In Table 12, 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 for state-level variation in real estate price appreciation. In Table 13, we present regression results using the baseline specification of equation (4) but only including the one-third of our sample who are least likely to own a home, judging by the proportion of housing units that are owner-occupied in their residential zip code, as measured by the 2000 U.S. Census.²³

In both tables, all of our MPC point estimates are negative, 5 of the 8 significantly so, and there is little evidence of attenuation of the MPCs relative to the baseline estimations. 6 of the 8 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.

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, especially for the young. Our results suggest that many investors are influenced by a reinforcement learning heuristic that leads high capital gains to encourage saving and discourage consumption, even when those capital gains are not useful for predicting future returns. This reinforcement learning framework reconciles our results with past work that has estimated positive MPCs out of cash windfalls and aggregate stock market movements.

²³ The average owner-occupied housing fraction (weighted by employees living in the zip code) in the lower third of the sample is 61.1%, versus 74.8% for the entire sample.

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.
- Carroll, Christopher D., 1992. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 1992(2), pp. 61-156.
- , 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112, pp. 1-57.
- 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, 2002. "Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance." In James M. Poterba, editor, *Tax Policy and the Economy*, Volume 16. Cambridge, MA: MIT Press, pp. 67-113.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2003. "Employees' Investment Decisions About Company Stock." In Olivia S. Mitchell and Stephen P. Utkus, editors, *Pension Design and Decision-Making Under Uncertainty*, forthcoming.
- Choi, James, David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2004. "For Better or For Worse: Default Effects and 401(k) Savings Behavior." In David A. Wise, editor, *Perspectives in the Economics of Aging*. Chicago, IL: University of Chicago Press.
- Erev, Ido, and Alvin E. Roth, 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88, pp. 848-881.
- Goetzmann, William, and Massimo Massa, 2004. "Daily Momentum and Contrarian Behavior of Index Fund Investors," *Journal of Financial and Quantitative Analysis*. Forthcoming.
- Graham, John R., and Campbell R. Harvey, 2003. "Expectations of Equity Risk Premia, Volatility, and Asymmetry." Mimeo, Duke University.

- 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.
- Honoré, Bo E, 1992. "Trimmed LAD and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects." *Econometrica* 60, pp. 533-565.
- Hurd, Michael D. and Michael J. Boskin, 1984. "The Effect of Social Security on Retirement in the Early 1970s." *Quarterly Journal of Economics* 99, pp. 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.
- Kearney, Melissa S. and Jenny Liao, 2004. "Lottery Players: What Are They Thinking?" Mimeo, Wellesley College.
- 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.
- Madrian, Brigitte C. and Dennis F. Shea, 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116, pp. 1149-1525.
- Patel, Jayendu, Darryll Hendricks, and Richard Zeckhauser, 1994. "Investment Flows and Performance: Evidence from Mutual Funds, Cross-Border Investments, and New Issues." In Ryuzo Sato, Richard M. Levich, and Rama Ramachandran, editors, *Japan, Europe, and International Financial Markets: Analytical and Empirical Perspectives*. New York: Cambridge University Press, pp. 51-72.
- Profit Sharing/401(k) Council of America, 2001. *44th Annual Survey of Profit Sharing and 401(k) Plans*. Chicago, IL: The Council.
- Roth, Alvin E., and Ido Erev, 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior* 8, pp. 164-212.

- Samuelson, William, and Richard Zeckhauser, 1988. "Status Quo Bias in Decision-Making." *Journal of Risk and Uncertainty* 1, 7-59.
- Sevak, Purvi, 2003. "Wealth Shocks and Retirement Timing: Evidence from the Nineties." Working Paper, Hunter College.
- 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.
- Vissing-Jorgensen, Annette, 2003. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear With Wealth? Evidence From Expectations and Actions." *NBER Macroeconomics Annual 2003*, forthcoming.

Table 1. Company Descriptions

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%
401(k) participation rate	80%	61%	86%	85%	83%
Average 401(k) balance	\$80,740	\$19,501	\$81,122	\$88,033	\$60,426
Maximum contribution rate (% of salary)	10% before-tax, 14% after-tax, 14% combined	15% before-tax	20% before-tax	25% before-tax and after-tax combined	1998-99: 14% before-tax 2000: 16% before-tax
Employer match	25% to 100% (varies by location) of first 6% of pay	25% of first 3% of pay	None until 2000, then 100% of first 1% of pay, 50% of next 4% of pay	50% of first 7% or 8% of pay (depends on union membership)	100% of first 3% of pay, 50% of next 3% of pay
Investment funds	1998: 3 bond, 3 large-cap, 1 mid-cap, 1 small-cap, 3 overseas, company stock. 1999: Added 1 bond, 1 large-cap, 1 overseas. 2000: Added 1 overseas and self-directed window.	1 cash, 1 bond, 3 pre-mix, 2 large-cap, 1 small-cap, 1 overseas, company stock	1 cash, 3 bond, 4 pre-mix, 8 large-cap, 5 mid-cap, 3 small-cap, 8 overseas, 3 sector, company stock	1998: 1 cash, 1 bond, 3 pre-mix, 1 large-cap, 1 mid-cap, 1 overseas, company stock 1999: Added 1 small cap, self-directed window	1 bond, 3 pre-mix, 5 large-cap, 1 small-cap, 1 overseas
Number of outstanding loans allowed	1 home loan, 1 general purpose loan	1	2	2	2
Hardship withdrawals allowed	Yes	Yes	Yes	Yes	Yes
Non-hardship withdrawal rules before age 59½	1 withdrawal allowed per month from after-tax, rollover, vested company match, and profit-share balances	After-tax and vested employer contribution money from grandfathered plans can be withdrawn at any time	Not allowed	After-tax and vested 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 the 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 pay cycle of December, normalized by income in the last pay cycle of the prior year (see equation (4)). 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 (Table 2), 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. Serial correlation is computed only over those participants who were active in the plan from the beginning of 1998 to the end of 2000.

	Normalized consumption growth	Normalized typical return wealth shock	Normalized adaptive wealth shock
Maximum	0.4251	5.2560	7.4919
99 th percentile	0.2455	0.4997	1.2310
90 th percentile	0.1292	0.0886	0.2111
75 th percentile	0.0630	0.0171	0.0418
50 th percentile	0.0097	-0.0023	-0.0004
25 th percentile	-0.0163	-0.0320	-0.0834
10 th percentile	-0.0579	-0.1192	-0.2928
1 st percentile	-0.1555	-0.6910	-1.1474
Minimum	-0.2864	-4.7744	-8.7422
Mean	0.0244	-0.0144	-0.0150
Std. Dev.	0.0774	0.1927	0.3846
Serial correlation	0.1027	-0.5102	-0.8735

Table 4. Regression of Consumption Growth on Normalized Wealth Shock

The dependent variable is year-over-year change in consumption in the last pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0037* (0.0016)	-0.0068** (0.0010)	-0.0050* (0.0024)	-0.0182** (0.0018)
$Shock_{t-1}/Y_{t-1}$	-0.0053** (0.0011)	0.0105** (0.0015)	-0.0011 (0.0020)	-0.0185** (0.0032)
Age_{t-1}	-0.0016** (0.0003)	-0.0015** (0.0003)	--	--
$Age^2_{t-1}/1000$	0.0092** (0.0034)	0.0082* (0.0034)	-0.1136** (0.0335)	-0.1552** (0.0336)
$\text{Log}(Y_{t-1})$	0.0045** (0.0011)	0.0035** (0.0011)	-0.2101** (0.0041)	-0.2106** (0.0041)
$\text{Log}(Tenure_{t-1})$	-0.0098** (0.0005)	-0.0104** (0.0005)	0.0388** (0.0022)	0.0403** (0.0022)
σ	0.0777** (0.0002)	0.0777** (0.0002)	--	--
N	69,581	69,581	54,054	54,054

* Significant at the 5% level

** Significant at the 1% level

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

The dependent variable is year-over-year change in consumption in the last pay cycle of December, normalized by salary in the last pay cycle 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 $t - 1$ 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t \times (Age < 30) / Y_{t-1}$	-0.0304 (0.0220)	-0.0267* (0.0109)	-0.1133** (0.0302)	-0.0929** (0.0172)
$Shock_t \times (30 \leq Age < 40) / Y_{t-1}$	-0.0072 (0.0048)	-0.0170** (0.0029)	-0.0175* (0.0073)	-0.0377** (0.0044)
$Shock_t \times (40 \leq Age < 50) / Y_{t-1}$	-0.0059* (0.0025)	-0.0082** (0.0015)	-0.0059 (0.0036)	-0.0186** (0.0024)
$Shock_t \times (50 \leq Age < 60) / Y_{t-1}$	-0.0002 (0.0025)	-0.0041** (0.0014)	0.0011 (0.0036)	-0.0104** (0.0029)
$Shock_t \times (Age \geq 60) / Y_{t-1}$	-0.0038 (0.0086)	-0.0075 (0.0053)	-0.0087 (0.0147)	-0.0054 (0.0112)
$Shock_{t-1} \times (Age < 30) / Y_{t-1}$	-0.0246 (0.0158)	0.0527* (0.0205)	-0.0603* (0.0297)	-0.1155** (0.0347)
$Shock_{t-1} \times (30 \leq Age < 40) / Y_{t-1}$	0.0021 (0.0033)	0.0155** (0.0048)	-0.0071 (0.0073)	-0.0356** (0.0088)
$Shock_{t-1} \times (40 \leq Age < 50) / Y_{t-1}$	-0.0033 (0.0017)	0.0123** (0.0022)	-0.0017 (0.0031)	-0.0167** (0.0044)
$Shock_{t-1} \times (50 \leq Age < 60) / Y_{t-1}$	-0.0064** (0.0015)	0.0094** (0.0020)	0.0017 (0.0027)	-0.0108* (0.0050)
$Shock_{t-1} \times (Age \geq 60) / Y_{t-1}$	-0.0112** (0.0029)	0.0108 (0.0064)	-0.0068 (0.0052)	-0.0037 (0.0168)
σ	0.0777** (0.0002)	0.0777** (0.0002)	--	--
N	69,581	69,581	54,054	54,054

* Significant at the 5% level

** Significant at the 1% level

Table 6. 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 $t - 1$. Estimates for company dummies are omitted. The columns labeled "Coefficient" present coefficient estimates from the probit. The columns labeled "Slope" present marginal effects evaluated at the means of the explanatory variables. Standard errors are in parentheses.

	Typical return wealth shock		Adaptive wealth shock	
	Coefficient	Slope	Coefficient	Slope
$Shock_t/Y_{t-1}$	0.3469* (0.1378)	0.1384** (0.0537)	0.0706 (0.0911)	0.0282 (0.0362)
$Shock_{t-1}/Y_{t-1}$	-0.0242 (0.0464)	-0.0096 (0.0184)	0.3293 (0.2412)	0.1314 (0.0956)
Age_{t-1}	0.0138 (0.0143)	0.0055 (0.0056)	0.0129 (0.0143)	0.0051 (0.0056)
$\text{Log}(Y_{t-1})$	-0.3147** (0.1196)	-0.1255** (0.0479)	-0.3318** (0.1207)	-0.1323** (0.0483)
$\text{Log}(Tenure_{t-1})$	0.1620* (0.0670)	0.0646* (0.0269)	0.1517* (0.0667)	0.0605 (0.0268)
N	1,457		1,457	

* Significant at the 5% level

** Significant at the 1% level

Table 7. 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 pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0011 (0.0066)	-0.0001 (0.0041)	-0.0240** (0.0092)	-0.0157** (0.0054)
$Shock_{t-1}/Y_{t-1}$	-0.0040 (0.0051)	0.0189** (0.0058)	-0.0081 (0.0083)	0.0023 (0.0096)
Age_{t-1}	-0.0022** (0.0008)	-0.0022** (0.0008)	--	--
$Age_{t-1}^2/1000$	0.0214* (0.0088)	0.0214* (0.0088)	-0.0486 (0.0578)	-0.0765 (0.0580)
$\text{Log}(Y_{t-1})$	-0.0071* (0.0028)	-0.0082** (0.0029)	-0.3208** (0.0111)	-0.3205** (0.0111)
$\text{Log}(Tenure_{t-1})$	-0.0136** (0.0014)	-0.0139** (0.0014)	0.0329** (0.0043)	0.0341** (0.0043)
σ	0.0787** (0.0006)	0.0786** (0.0006)	--	--
N	11,101	11,101	8,882	8,882

* Significant at the 5% level

** Significant at the 1% level

Table 8. 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 pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0031 (0.0016)	-0.0061** (0.0011)	-0.0061* (0.0025)	-0.0182** (0.0019)
$Shock_{t-1}/Y_{t-1}$	-0.0054** (0.0011)	0.0104** (0.0015)	-0.0019 (0.0020)	-0.0188** (0.0032)
Age_{t-1}	-0.0013** (0.0003)	-0.0012** (0.0003)	--	--
$Age^2_{t-1}/1000$	0.0064 (0.0035)	0.0054 (0.0035)	-0.1113** (0.0353)	-0.1475** (0.0353)
$\text{Log}(Y_{t-1})$	0.0033** (0.0011)	0.0022* (0.0011)	-0.2052** (0.0043)	-0.2056** (0.0043)
$\text{Log}(Tenure_{t-1})$	-0.0100** (0.0005)	-0.0106** (0.0005)	0.0383** (0.0024)	0.0399** (0.0024)
σ	0.0765** (0.0002)	0.0764** (0.0002)	--	--
N	64,668	64,668	47,474	47,474

* Significant at the 5% level

** Significant at the 1% level

Table 9. Regression of Consumption Growth on Normalized Wealth Shock for Unmarried Participants

The dependent variable is year-over-year change in consumption in the last pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0062 (0.0037)	-0.0118** (0.0024)	-0.0097* (0.0048)	-0.0229** (0.0034)
$Shock_{t-1}/Y_{t-1}$	-0.0054** (0.0019)	0.0168** (0.0034)	-0.0011 (0.0032)	-0.0212** (0.0064)
Age_{t-1}	-0.0013* (0.0005)	-0.0012* (0.0005)	--	--
$Age_{t-1}^2/1000$	0.0067 (0.0057)	0.0053 (0.0057)	-0.1368* (0.0630)	-0.1948** (0.0629)
$\text{Log}(Y_{t-1})$	-0.0058** (0.0021)	-0.0072** (0.0021)	-0.2375** (0.0075)	-0.2375** (0.0075)
$\text{Log}(Tenure_{t-1})$	-0.0096** (0.0009)	-0.0105** (0.0009)	0.0411** (0.0036)	0.0427** (0.0036)
σ	0.0833** (0.0004)	0.0832** (0.0004)	--	--
N	26,879	26,879	17,792	17,792

* Significant at the 5% level

** Significant at the 1% level

Table 10. 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 consumption in the last pay cycle of December, normalized by salary in the last pay cycle of 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0000 (0.0016)	-0.0015 (0.0010)	0.0004 (0.0023)	-0.0165** (0.0017)
$Shock_{t-1}/Y_{t-1}$	0.0015 (0.0010)	0.0059** (0.0015)	0.0004 (0.0018)	-0.0150** (0.0030)
Age_{t-1}	-0.0014** (0.0003)	-0.0014** (0.0003)	--	--
$Age^2_{t-1}/1000$	0.0080* (0.0033)	0.0080* (0.0032)	-0.2901** (0.0333)	-0.3079** (0.0332)
$\text{Log}(Y_{t-1})$	-0.0017 (0.0010)	-0.0023* (0.0010)	-0.9486** (0.0082)	-0.9474** (0.0081)
$\text{Log}(Tenure_{t-1})$	-0.0107** (0.0005)	-0.0109** (0.0005)	0.0559** (0.0059)	0.0687** (0.0060)
σ	0.0768** (0.0002)	0.0768** (0.0002)	--	--
N	68,190	68,190	52,278	52,278

* Significant at the 5% level

** Significant at the 1% level

**Table 11. 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0008 (0.0015)	-0.0016 (0.0009)	0.0004 (0.0019)	-0.0152** (0.0014)
$Shock_{t-1}/Y_{t-1}$	0.0023* (0.0009)	0.0062** (0.0013)	0.0015 (0.0015)	-0.0116** (0.0025)
Age_{t-1}	-0.0011** (0.0003)	-0.0011** (0.0003)	--	--
$Age^2_{t-1}/1000$	0.0048 (0.0030)	0.005 (0.0030)	-0.3064** (0.0276)	-0.3258** (0.0275)
$\text{Log}(Y_{t-1})$	0.0001 (0.0009)	-0.0005 (0.0010)	-0.9762** (0.0067)	-0.9753** (0.0067)
$\text{Log}(Tenure_{t-1})$	-0.0105** (0.0005)	-0.0107** (0.0005)	0.0579** (0.0050)	0.0709** (0.0050)
σ	0.0715** (0.0002)	0.0715** (0.0002)	--	--
N	69,581	69,581	54,054	54,054

* Significant at the 5% level

** Significant at the 1% level

**Table 12. Regression of Consumption Change on Normalized Wealth Shock
with State \times Year Effects**

The dependent variable is year-over-year change in consumption in the last pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0035* (0.0016)	-0.0071** (0.0010)	-0.0040 (0.0024)	-0.0157** (0.0018)
$Shock_{t-1}/Y_{t-1}$	-0.0056** (0.0011)	0.0100** (0.0015)	-0.0013 (0.0020)	-0.0149** (0.0032)
Age_{t-1}	-0.0014** (0.0003)	-0.0013** (0.0003)	--	--
$Age^2_{t-1}/1000$	0.0075* (0.0034)	0.0065 (0.0034)	-0.1360** (0.0338)	-0.1725** (0.0338)
$\text{Log}(Y_{t-1})$	0.0024* (0.0011)	0.0015 (0.0011)	-0.2017** (0.0041)	-0.2024** (0.0041)
$\text{Log}(Tenure_{t-1})$	-0.0077** (0.0005)	-0.0083** (0.0005)	0.0386** (0.0022)	0.0400** (0.0022)
σ	0.0767** (0.0002)	0.0766** (0.0002)	--	--
N	69,581	69,581	54,054	54,054

* Significant at the 5% level

** Significant at the 1% level

Table 13. Regression of Consumption Growth on Normalized Wealth Shock for Participants in the Lowest Third of the Home Ownership Probability Distribution

The dependent variable is year-over-year change in consumption in the last pay cycle of December, normalized by salary in the last pay cycle 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. Standard errors are in parentheses.

	Tobit		Employee fixed effects	
	Typical return wealth shock	Adaptive wealth shock	Typical return wealth shock	Adaptive wealth shock
$Shock_t/Y_{t-1}$	-0.0027 (0.0031)	-0.0066** (0.0020)	-0.0065 (0.0044)	-0.0180** (0.0032)
$Shock_{t-1}/Y_{t-1}$	-0.0068** (0.0019)	0.0153** (0.0028)	0.0006 (0.0033)	-0.0160** (0.0055)
Age_{t-1}	-0.0018** (0.0005)	-0.0017** (0.0005)	--	--
$Age_{t-1}^2/1000$	0.0123* (0.0059)	0.0114 (0.0059)	-0.0572 (0.0532)	-0.0946 (0.0532)
$\text{Log}(Y_{t-1})$	-0.0009 (0.0019)	-0.0024 (0.0019)	-0.2562** (0.0076)	-0.2561** (0.0076)
$\text{Log}(Tenure_{t-1})$	-0.0081** (0.0009)	-0.0089** (0.0009)	0.0424** (0.0038)	0.0439** (0.0038)
σ	0.0817** (0.0004)	0.0817** (0.0004)	--	--
N	23,366	23,366	17,392	17,392

* Significant at the 5% level

** Significant at the 1% level