

Reinforcement Learning and Investor Behavior^{*}

James J. Choi
Yale University and NBER

David Laibson
Harvard University and NBER

Brigitte C. Madrian
Harvard University and NBER

Andrew Metrick
University of Pennsylvania and NBER

Current draft: September 14, 2007

Abstract: What affects individual investors' willingness to invest in an asset? This paper presents evidence that — when there is no salient reference purchase price — investors tend to be *return chasers* and *variance avoiders* with respect to their *idiosyncratic* history with the asset. Using administrative panel data on 25,000 401(k) accounts at five firms, we find that an investor's 401(k) contribution rate increases more if she has recently experienced a higher 401(k) portfolio return and/or a lower 401(k) return variance. We find no evidence that this behavior is welfare-improving. These results are explained by a naïve reinforcement learning heuristic: investors expect that investments in which they experienced past success will be successful in the future, whether or not such a belief is logically justified. Consistent with reinforcement learning's Power Law of Practice, return chasing and variance avoidance diminish with age.

^{*} This paper is a significantly revised version of a draft that circulated under the title "Consumption-Wealth Comovement of the Wrong Sign." 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 the comments of Thomas Knox, Matthew Shapiro, Bill Zame, and seminar participants at Berkeley, Harvard, the NBER Monetary Economics Program Meeting, UCLA, and RAND. We are also indebted to John Beshears, Ian Dew-Becker, Keith Ericson, Fuad Faridi, John Friedman, Chris Nosko, 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.

What affects individual investors' willingness to invest in an asset? A large literature on the "disposition effect" has shown that their idiosyncratic history with that asset plays a significant role. Investors are reluctant to sell assets that have fallen below their purchase price and more likely to sell assets that have risen above their purchase price.¹ This behavior is anomalous because the asset's purchase price is investor-specific and already sunk, and hence should not affect the selling decision in the absence of capital gains taxes.² Odean (1998) shows that investors are also more likely to buy additional shares of stocks they already own if they have unrealized losses in those stocks. The most common explanation for the disposition effect is that prospect theory preferences (Kahneman and Tversky, 1979) cause investors to experience disutility from making a sale below the "reference price" at which they bought the asset, and to be risk-seeking for assets that are mentally classified in the loss domain.³

In this paper, we study how idiosyncratic history matters when there is *no* salient reference purchase price. Such situations commonly arise when periodic asset purchases are automatically made for an investor over time. This causes the investor's holdings to have many different purchase prices for each asset, and purchases usually have occurred without the investor's direct oversight. Therefore, it is difficult for the investor to mentally establish a single reference price below which his investment is in the loss domain. Automatic purchases occur most frequently in employer-sponsored 401(k) plans⁴—the setting we focus on—which divert a pre-specified proportion of each paycheck to invest in the plan assets until the employee elects

¹ The seminal papers are Shefrin and Statman (1985) and Odean (1998). The disposition effect has since been documented by many other researchers, including Grinblatt and Keloharju (2001), Shapira and Venezia (2001), Genesove and Mayer (2001), Wermers (2003), Coval and Shumway (2005), Locke and Mann (2005), and Dhar and Zhu (2006).

² Introducing capital gains taxes should make investors more prone to sell losers (Constantinides, 1984), which is the opposite of what they actually do.

³ See, however, Barberis and Xiong (2006) for an argument that prospect theory does not robustly produce the disposition effect.

otherwise or leaves the firm. Many mutual fund companies also allow investors to set up a recurring transfer from their bank account to the fund company to buy fund shares.

Using administrative panel data on approximately 25,000 individual 401(k) accounts at five companies, we find that in the absence of a salient reference price, investors act as *return chasers* and *variance avoiders* with respect to their idiosyncratic 401(k) returns history. A one standard deviation increase in an investor's rate of return in her overall 401(k) portfolio relative to her coworkers during year t increases the change in her 401(k) contribution rate between year-ends $t - 1$ and t by 0.13 percentage points of income. Conversely, a one standard deviation increase in the variance of an investor's 401(k) return relative to her coworkers during year t lowers her (cumulative) contribution rate change between year-ends $t - 1$ and $t + 1$ by 0.35 percentage points. These effects are economically significant. By comparison, the average annual contribution rate change in our sample is 0.30 percentage points.

Our empirical methodology allows us to rule out many rational explanations for the return chasing and variance avoidance we observe. Because we include individual fixed effects in our regressions, we control for unobserved time-invariant heterogeneity that drives differential average 401(k) contribution rate changes. Because we include company \times year fixed effects, we hold constant public news about expected returns for the plan's investment options. Essentially, we are comparing *contemporaneous* contribution rate changes across individuals within the *same* 401(k) plan. Investor A should not believe the mutual funds in the 401(k) investment menu have a different expected return distribution than her coworker, Investor B, does simply because A's chosen investments in the plan outperformed B's chosen investments in the plan. Our results are also robust to controlling for wealth effects and year-specific shocks that are correlated with an

⁴ At year-end 2005, 401(k) plans in aggregate held \$2.4 trillion of retirement savings owned by 47 million individuals (Holden and VanDerhei, 2006).

individual's absolute (dollar) and relative (percent) 401(k) allocations to equities, bonds, and cash at year-end $t - 1$.

We investigate the possibility that return chasing and variance avoidance are the result of investors rationally learning about their skill at asset allocation within the 401(k). If an investor can infer from her high realized performance that she has 401(k) asset allocation skill, she could reasonably decide to increase her 401(k) contributions on that basis. But contrary to the rational learning hypothesis, we find no evidence that high past 401(k) alphas predict high future 401(k) alphas. If anything, a high alpha in the current year predicts a low alpha in the following year.

Our findings are explained by a naïve reinforcement learning heuristic: investors expect that investments in which they *personally* experienced past success will be successful in the future, whether or not such a belief is logically justified. A key feature of reinforcement learning models is the Power Law of Practice: learning curves are steep initially and then level out as the stock of reinforcements increases and each additional stimulus constitutes a smaller proportional addition to the stock (Roth and Erev, 1995). Consistent with the Power Law of Practice, we find that both return chasing and variance avoidance diminish with age. Nevertheless, our point estimates indicate that investors continue to follow these heuristics into their sixties.

The return chasing we observe contrasts sharply with the contrarian behavior associated with the disposition effect, where investors are prone to keep or increase their holdings of losers. The two phenomena can be reconciled by decomposing idiosyncratic history effects into a belief-based component and a preference-based component. Past return performance positively affects estimates of future return performance through reinforcement learning. In the absence of a reference price, these beliefs induce return chasing and variance avoidance. However, once a reference price becomes salient, the loss aversion induced by prospect theory preferences is

activated and dominates the reinforcement learning effect, leading to reluctance to close out losing positions and a propensity to increase risk-taking in those securities.⁵

Our results complement Barber, Odean, and Strahilevitz (2004), who document brokerage investors' propensity to repurchase individual stocks they previously sold for a gain while shunning individual stocks they previously sold for a loss. Barber, Odean and Strahilevitz find that purchased stocks previously sold for a gain do not subsequently underperform relative to benchmarks based on size and book-to-market. Therefore, conditional on making a purchase, the propensity to buy previously profitable stocks appears to be welfare-neutral. In our setting, however, welfare will generally be affected by changes in an employee's 401(k) contributions, which are tax-advantaged and often garner a matching employer contribution.

Our paper is also related to the large literature finding that investors chase mutual fund returns.⁶ Because these other papers identify return chasing using variation in mutual fund returns, it is difficult for them to tell whether return chasing is motivated by rational updating on truly useful information embedded in fund returns, or if investors are blindly chasing returns *per se*.⁷ We identify return chasing holding fixed mutual fund returns and using variation in returns

⁵ One might suspect that we are not finding contrarian behavior simply because 401(k) investments are mostly in mutual funds rather than individual stocks, not because there is no salient reference purchase price. However, Barber, Odean, and Zheng (2000) document the disposition effect in brokerage account mutual fund holdings. Such holdings are much more likely to have a salient reference purchase price, since they were probably not accumulated through an automatic purchase plan.

⁶ This literature includes Ippolito (1992), Goetzmann and Peles (1993), Patel, Zeckhauser, and Hendricks (1994), Sirri and Tufano (1998), Chevalier and Ellison (1997), and Goetzmann and Massa (2002).

⁷ The rationality of mutual fund returns-chasing has been widely debated. Gruber (1996) and Zheng (1999) find that returns-chasing is profitable, whereas Frazzini and Lamont (2006) come to the opposite conclusion. Sapp and Tiwari (2004) conclude that any profitability from individuals' returns-chasing can be explained by the stock momentum effect (Jegadeesh and Titman, 1993). Empirical studies of mutual fund performance persistence that do not specifically examine retail flows include Grinblatt and Titman (1992), Hendricks, Patel, and Zeckhauser (1993), Goetzmann and Ibbotson (1994), Brown and Goetzmann (1995), Elton, Gruber, and Blake (1996), and Carhart (1997). Berk and Green (2004) argue that returns-chasing without mutual fund performance persistence is to be expected when rational investors delegate funds to portfolio managers with diminishing returns to scale in security selection. In their model, money flows to a manager who has demonstrated stock-picking skill until the manager can no longer deliver superior returns going forward. Such diminishing returns, however, can not explain the key finding of our analysis: 401(k) investors are more prone to chase returns that they have *personally* experienced, rather than those they have observed but not experienced.

investors have *personally* experienced in their 401(k).⁸ If a naïve reinforcement learning heuristic causes 401(k) investors to become more willing to hold 401(k) assets when their 401(k) has idiosyncratically performed well, then similarly naïve motivations may drive the higher willingness to hold mutual funds that have recently performed well.

Finally, our paper adds to a small but growing empirical literature on investor learning. Other papers have studied how learning improves investing skill, as manifested in higher portfolio returns or decreasing strength of the disposition effect (Nicolosi, Peng, and Zhu, 2004; Feng and Seasholes, 2005; Seru, Shumway, and Stoffman, 2006). In contrast, our paper focuses on how investors update their portfolios in response to irrelevant information, although we do find that this responsiveness attenuates with experience

The rest of the paper proceeds as follows. Section I describes our 401(k) data. Although ours is the first study to use these specific records, other papers have exploited different sets of individual-level 401(k) data to study savings rates (Huberman, Iyengar, and Jiang, 2006), asset allocation (Benartzi and Thaler, 2001; Agnew, Balduzzi, and Sundén, 2003; Huberman and Jiang, 2006), trading (Choi, Laibson, and Metrick, 2002; Agnew, Balduzzi, and Sundén, 2003), and investment in employer stock (Huberman and Sengmueller, 2004; Choi, Laibson, Madrian, and Metrick, 2004). Section II explains the framework within which we conduct our empirical estimation. Section III presents our results, and Section IV considers alternative interpretations of the results. Section V concludes.

⁸ To our knowledge, the only other paper that uses a similar identification method is Calvet, Campbell, and Sodini (2007), who document that within a given time period, Swedish investors are more likely to stop holding risky assets altogether after experiencing poor mutual fund returns. The authors are primarily concerned with rebalancing behavior and do not connect poor portfolio performance with anything other than overall risky asset participation status.

I. Data description

Our data come from a large benefits record-keeping firm. We have panel data for five companies that start when our data provider became the plan administrator at each company and end at year-end 2000. These data contain the date, amount, and type of every transaction made in these firms' 401(k) plans by every participant. In addition, we have year-end cross-sectional snapshots from 1998, 1999, and 2000 for all active employees that include demographic information such as birth date, hire date, gender, compensation, marital status, and state of residence (e.g. California). The year-end cross-sections also contain point-in-time 401(k) information, including the contribution rate in effect during the final pay period of the year, total balances, and asset allocations.

Table I 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. Equally weighting each company, 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.

At each of these firms, employees can choose a contribution rate that is an integer percentage of their salary. The contribution rate determines how much of each paycheck is deducted and contributed to the plan, and it remains in effect until the employee actively changes

it. All of our companies offer matching contributions proportional to employee contributions up to a threshold, although Company C did not introduce its match until 2000. For example, employees who contributed at least 3% of their pay at Company B received an additional contribution from the company equal to 0.75% their pay.

The large majority of the plans' investment options are mutual funds. Every plan offers at least eight mutual funds, including at least one fixed-income fund. The most important investment option that is not a mutual fund is employer stock, which is offered by four of our five plans. In addition, Companies A and D added a self-directed window to their plans in 2000 and 1999, respectively. Self-directed windows allow participants to buy and sell individual stocks using their 401(k) balances. We do not observe transactions within the self-directed windows, although we do know the total balances held in the windows at each year-end. Among plan participants in Companies A and D, 1.1% and 8.0%, respectively, had any balances in the self-directed window at year-end 2000. Conditional upon having any money in the window, participants in Companies A and D held on average 34.1% and 28.1% of their 401(k) balances in the window, respectively.

All of the plans allow participants to take hardship withdrawals from and loans against their 401(k) plan balances, and only one does not allow non-hardship withdrawals. These provisions make 401(k) savings in the companies we study more liquid than for the typical 401(k) participant at the time.⁹ 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

⁹ 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.

Internet 24 hours a day. The direct transaction cost of changing one's contribution rate in these plans is therefore minimal.

II. Empirical methodology

Our empirical objective is to estimate the relationship between changes in an individual's 401(k) contribution rate and the first two moments of 401(k) returns. We define the one-year 401(k) rate of return, R , as the sum of the monthly percent returns of the 401(k) over a one-year period. More formally, let i index investors, f index 401(k) investment funds, t index years, and m index months within a year. Then

$$R_{i,t} = \sum_{m=1}^{12} \sum_f w_{i,f,t,m} r_{f,t,m}, \quad (1)$$

where $w_{i,f,t,m}$ is the fraction of i 's 401(k) portfolio held in fund f at the beginning of the m th month of year t , and $r_{f,t,m}$ is the return of fund f during the m th month of year t . We define $\sigma^2(R_{i,t})$ as the annualized variance of the twelve monthly returns that comprise $R_{i,t}$.

We adopt a flexible functional form for the determinants of an individual's 401(k) contribution rate. Let $C_{i,t}$ be the 401(k) contribution rate, measured as a percent of salary, in effect for individual i at year-end t .¹⁰ Then

$$C_{i,t} = g_i(\text{age}_{i,t}) + \beta_1 R_{i,t} + \beta_2 R_{i,t-1} + \beta_3 \sigma^2(R_{i,t}) + \beta_4 \sigma^2(R_{i,t-1}) + \mathbf{X}_{i,t} \boldsymbol{\gamma} + \varepsilon_{i,t}, \quad (2)$$

where $g_i(\cdot)$ is a function specific to investor i , $\text{age}_{i,t}$ is the investor's age, $\mathbf{X}_{i,t}$ is a vector of other control variables defined as of year-end t , and $\varepsilon_{i,t}$ is the residual term.¹¹ The function g_i could

¹⁰ By using the instantaneous contribution rate at year-end rather than the total contributions during the year, we can be sure that all the information in the explanatory variables was potentially available to the investor before she made her choice of the dependent variable.

¹¹ We could also model the contribution rate *change* as a function of the level of R and other controls. We continue to find significant returns-chasing under this alternative specification. We prefer the specification in the main text

vary across investors due to unobserved differences (e.g. discount rates, risk aversion, expected income growth, background risk) that alter the optimal solution to the lifecycle consumption-investment problem. By controlling for g , we control for year-over-year contribution changes that would have occurred due to lifecycle considerations, regardless of 401(k) returns. We include contemporaneous returns $R_{i,t}$ and their variance $\sigma^2(R_{i,t})$. We also include lagged returns and the variance of lagged returns, $R_{i,t-1}$ and $\sigma^2(R_{i,t-1})$, to allow for the possibility of a sluggish response to 401(k) performance.

We assume that the function $g_i(\text{age}_{i,t})$ is locally well-approximated by a first-order Taylor expansion around the investor's age at year-end 1999 (the middle year in our sample of contribution rates):

$$g_i(\text{age}_{i,t}) \approx g_i(\text{age}_{i,1999}) + \frac{\alpha_i}{2}(\text{age}_{i,t} - \text{age}_{i,1999}). \quad (3)$$

Substituting (3) into (2) and first-differencing yields an equation with an individual fixed effect in contribution rate changes:

$$\Delta C_{i,t} = \alpha_i + \beta_1 \Delta R_{i,t} + \beta_2 \Delta R_{i,t-1} + \beta_3 \Delta \sigma^2(R_{i,t}) + \beta_4 \Delta \sigma^2(R_{i,t-1}) + \Delta \mathbf{X}_{i,t} \boldsymbol{\gamma} + \Delta \varepsilon_{i,t}. \quad (4)$$

We estimate (4) using least-squares regression, first-differencing the equation to eliminate the individual fixed effect. Hence, the dependent variable is ultimately the second-difference of $C_{i,t}$, or the change in contribution rate changes. Note that because we have only three years of contribution rate data, there is only one second-differenced C observation per person. We cluster our standard errors at the company \times state level in case peer effects or information spillovers

because making the contribution rate change a function of the return level allows the contribution rate level to drift without bound in theory.

cause dependence in contribution rate changes between coworkers in the same office (Duflo and Saez, 2003; Hong, Kubik, and Stein, 2004; Ivkovich and Weisbenner, 2007).¹²

The presence of the individual fixed effect imposes the requirement that all employees in our regressions have two contribution rate change observations. We also need four full years of capital gains data in order to estimate the coefficients on both contemporaneous and lagged ΔR . Thus, our sample is limited to workers who have been actively employed at a sample firm and continuously enrolled in the 401(k) plan from January 1, 1997 to December 31, 2000. Company E's data start on March 31, 1997, when our data provider assumed administrative services for its plan, so we instead require its workers to be actively employed and continuously enrolled in the plan from March 31, 1997 to December 31, 2000.¹³

Contributions to 401(k) plans are usually made with before-tax money. However, some of our sample plans allow contributions using after-tax money as well. We add the before-tax and (if the plan offers the option) after-tax 401(k) contribution rates in effect for the last pay period of 1998, 1999, or 2000 to calculate $C_{i,t}$ in each of these years. We include employees whose contribution rate is zero, provided that they had a positive contribution rate or balance at some time in the past.¹⁴ We also require that individuals have salaries greater than \$20,000 in 1998 because a large fraction of those with salaries under \$20,000 are part-time employees who are likely to direct less attention to the 401(k) than full-time employees.¹⁵ In addition, we trim workers who have a one-year income growth observation greater than 30% or less than -20%,

¹² Consistent with there being only weak geographic effects in contribution rates, our standard errors are barely affected by clustering relative to assuming that all observations are independent. In contrast, the standard errors in our portfolio return persistence analysis, presented in Section V.A, are greatly increased by clustering.

¹³ We assign a zero 401(k) return to Company E employees for the first three months of 1997. Our results are qualitatively similar if we drop Company E from the sample.

¹⁴ Employees with no balances in the plan are excluded from our analysis because our key explanatory variable, the individual's rate of return on plan assets, is only defined for those with assets in the plan.

which roughly corresponds to removing the top 2% and bottom 2% of the income growth distribution. These deleted outliers are likely caused by transitions between part-time and full-time work status.

Finally, we drop individuals if their 1998 salary is high enough that, by contributing at the plan's maximum before-tax contribution rate, they could exceed the \$10,000 statutory limit on 1998 before-tax 401(k) contributions. The reason we impose this selection rule is that a highly paid employee could contribute enough that he hits the before-tax dollar limit midway through the year. For the rest of the year, his before-tax contribution rate is frozen at 0 and does not reflect any desired change.¹⁶

All of our specifications include the log of the employee's tenure at the company and company dummies interacted with year dummies in the $\mathbf{X}_{i,t}$ vector. The company \times year dummies control for public news about future asset returns that affect optimal contribution rates, as well as news that is specifically relevant to employees of each company.¹⁷ In many specifications, we also control for wealth effects by adding contemporaneous and lagged 401(k) capital gains normalized by current income, $CapitalGain_{i,t}/Y_{i,t}$ and $CapitalGain_{i,t-1}/Y_{i,t}$, where $CapitalGain_{i,t}$ is investor i 's 401(k) dollar capital gain during year t and $Y_{i,t}$ is the investor's annual salary. We calculate $CapitalGain_{i,t}$ by taking the difference in balances between year-end t and $t - 1$ and then subtracting contributions, rollovers into the plan, and loan repayments during year t and adding back withdrawals and new loans during year t . We normalize $CapitalGain$ (a

¹⁵ 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.

¹⁶ We also drop a small number of Company A employees who are eligible to contribute to the company's deferred compensation plan.

¹⁷ The company \times year dummies reduce to a single company dummy in the final double-differenced equation (4).

variable whose unit is dollars) by income because the dependent variable in our regressions (contribution rate) is also expressed as a percent of income.

Unfortunately, we cannot calculate portfolio returns, R , including returns in the self-directed windows at Companies A and D, since we do not observe monthly window balances. The two capital gains variables, however, do include dollar gains realized in the window. Our contribution change results are robust to excluding Companies A and D from the sample.

III. Results

A. Summary statistics

The selection criteria described in Section II leave us with 49,248 contribution rate change observations on 24,624 employees. Table II reports summary statistics for contribution rate changes and our portfolio return variables. From 1998 to 2000, the median annual contribution rate change is zero, and the mean change is 0.30 percentage points of income. Between 1998 and 1999, 20.6% of our sample investors changed their contribution rate, and 22.4% changed their contribution rate between 1999 and 2000 (these specific numbers are not reported in the table). Over the two years, 35.1% of investors made at least one contribution rate change.

Pooled across 1997 to 2000, the one-year 401(k) rate of return, R , has a median of 9.9% and a mean of 11.9%. Reflecting the dramatic late-1990s bull market and subsequent crash, R has a wide distribution; its pooled cross-sectional standard deviation is 16.9%. The volatility of portfolio returns, $\sigma^2(R)$, also exhibits wide variation across individuals due to differing portfolio shares allocated to equities and particularly to employer stock. The typical volatility is quite high, since many plan participants held significant amounts of their employer's stock. Our

companies' monthly stock returns generally experienced annualized standard deviations well over 100% during the sample period. The dollar capital gain normalized by income, $CapitalGain/Y$, has an economically narrower range because most investors' 401(k) balances are modest compared to their income. The mean and median of $CapitalGain/Y$ are 0.04 and 0.09, respectively, and its standard deviation is 0.30.

The extreme two percentiles of R and $CapitalGain/Y$ contain especially large positive and negative returns. The minimum and maximum values of R are -82.6% and 127.6% , and the corresponding values for $CapitalGain/Y$ are -8.49 and 7.87 . The R outliers are due to an OTC fund, a technology fund, two growth funds, two international funds, and one company's stock; the $CapitalGain/Y$ outliers are investors with large 401(k) balances who invested in funds with extreme returns. The estimates that we report below are not driven by these outliers. We have re-run our regressions trimming the extreme two percentiles of R , $\sigma^2(R)$, or $CapitalGain/Y$ and find that our results are nearly unchanged.

B. Main contribution rate change regressions

Table III presents the coefficients from estimating equation (4). The first column shows estimates from the baseline specification, which includes first-differenced contemporaneous and lagged 401(k) return and volatility, log tenure, and company-year dummies as explanatory variables. We find that a one standard deviation increase in year t return causes the change in the 401(k) contribution rate between year-ends $t - 1$ and t to rise by $16.87 \times 0.0078 = 0.13$ percentage points of income, an effect that is significant at the 1% level. There is no effect in year $t + 1$.

In contrast, a one standard deviation increase in year t volatility causes the change in the contribution rate between year-ends $t - 1$ and t to fall by $6,307/1000 \times 0.029 = 0.18$ percentage points in year t . The contribution rate change continues to fall an additional 0.17 percentage points between year-ends t and $t + 1$. Both the contemporaneous and lagged variance-avoidance effects are significant at the 1% level.

To assess the economic significance of these effects, recall that the average annual contribution rate increase is 0.30 percentage points of income. Thus, the 0.13 percentage point effect of 401(k) returns and the 0.35 percentage point two-year effect of volatility are substantial relative to the mean.

Because we are including company \times year dummies in our regression, we are controlling for public news about expected asset returns and news specifically relevant to employees of each company. Holding fixed news, an investor should not update his beliefs about the future returns of his 401(k)'s investment options differently based upon the fact that his portfolio did better or worse than his coworkers' portfolio over the same time period. Yet we find that employees do invest more in their 401(k) when their portfolio performance was better than their coworkers'. Because we are including individual fixed effects in our regression, we are also controlling for sources of time-invariant investor heterogeneity—such as risk aversion, time preference, human capital, etc.—that may affect changes in contribution rates.

Even after controlling for time shocks common to each company, one might worry that there is cross-sectional heterogeneity in the time shock that is correlated with 401(k) portfolio returns. One candidate for such a correlated shock is the wealth effect from the 401(k) portfolio capital gain itself. However, the second column of Table III shows that controlling for the

contemporaneous and lagged normalized dollar capital gains in the 401(k) barely affects the coefficients on return and volatility.

Another potential concern is that a series of economic news arrived during our sample period that differentially affected the type of people who tend to hold relatively more equities (e.g. news about the return to high-skill human capital). Because asset class allocations are in turn correlated with portfolio returns, this could confound our identification. To account for this possibility, the third column of Table III adds interactions between year dummies and three variables: the dollar amount of the individual's portfolio held in equities, bonds, and cash, all as a *fraction of income* at the prior year-end. In the last column of Table III we also add interactions between year dummies and two variables: the *fraction of one's 401(k)* allocated to equities and to bonds (also at the prior year-end). Even with these additional controls, we continue to estimate large and statistically significant return chasing and variance avoidance, and the point estimates remain similar to those in the baseline specification of column 1. In the most comprehensive specification, we find that a one standard deviation increase in portfolio returns in year t increases the 401(k) contribution rate change in year t by 0.12 percentage points, and a one standard deviation increase in the volatility of returns in year t decreases the 401(k) contribution rate change by 0.16 percentage points in year t and 0.17 percentage points in year $t + 1$.

These findings are consistent with individual investors following a naïve reinforcement learning heuristic: investors expect that investments in which they *personally* experienced past rewards will be rewarding in the future, whether or not such a belief is logically justified. Reinforcement learning models have had success in predicting subject choices in experiments (Roth and Erev, 1995; Erev and Roth, 1998). Reinforcement learning is often a sensible heuristic because future rewards *are* positively correlated with recent rewards in many domains. Financial

markets are a rare exception. Charness and Levin (2003) show that when an (optimal) Bayesian updating rule conflicts with a reinforcement learning rule, experimental subjects' choices shift towards the erroneous option that reinforcement learning recommends.

C. Interactions with age and salary

A key feature of reinforcement learning models is the Power Law of Practice: learning curves are steep initially and then level out as the stock of reinforcements increases and each additional stimulus constitutes a smaller proportional addition to the stock (Roth and Erev, 1995). Reinforcement learning therefore predicts that the contribution rate of young investors is more responsive to idiosyncratic portfolio performance than that of old investors.

The regression in the first column of Table IV tests this prediction by interacting contemporaneous and lagged change in 401(k) return and volatility, ΔR and $\Delta\sigma^2(R)$, with de-meaned investor age at year-end 1998. For brevity, we show only the most comprehensive regression specification that controls for contemporaneous and lagged normalized *CapitalGain*, asset class balance \times year dummies, and asset class portfolio share \times year dummies. We indeed find that both return chasing and variance avoidance attenuate with age. The age interaction with the contemporaneous change in 401(k) return is only significant at the 10% level ($p = 0.09$). However, simply removing the statistically and economically insignificant age interaction with the lagged change in returns from the regression causes the age interaction with contemporaneous returns to become significant at the 1% level while retaining the same point estimate. We have greater statistical power for the age interactions with contemporaneous and lagged volatility changes, which are both significant at the 1% level and of similar magnitudes.

Even though responsiveness to idiosyncratic portfolio returns decreases with age, investors nonetheless exhibit reinforcement learning behavior for most of their lives. The point estimates indicate that return chasing continues until age 68.¹⁸ Variance avoidance diminishes more swiftly, but both contemporaneous and lagged variance-avoidance persists through age 54.

One might suspect that higher-income investors would be less prone to naïve reinforcement learning, since income is a proxy for financial sophistication. The second column of Table IV examines whether this is the case by interacting contemporaneous and lagged 401(k) return change and volatility change with 1998 log salary. Surprisingly, income has no significant attenuating effect on reinforcement learning tendencies, at least within the low-to-moderate income investor population in our regressions.

The final column of Table IV interacts return changes and with both age and log income. We see that the conclusions drawn from the first two columns are robust to allowing this simultaneous interaction. Age continues to attenuate the force of reinforcement learning, whereas income does not, and the point estimates of the interactions are nearly identical to those in the first two columns.

IV. Alternative explanations

We now consider alternative mechanisms that could generate the return-chasing and variance-avoidance results presented above.

¹⁸ Age is de-meaned in the Table 4 to facilitate interpretation of the uninteracted ΔR and $\Delta\sigma^2(R)$ coefficients. The mean age at year-end 1998 in the regression sample is 43.4. Therefore, return chasing drops to zero at age $43.4 + 0.0059/0.0024 \times 10 = 68.0$ in column 1.

A. Learning about investing skill

Investors who experience high 401(k) returns with low variance may be learning that they have greater skill at 401(k) asset allocation than their coworkers who experience low 401(k) returns with high variance. Therefore, it may be rational for investors with better performance to allocate more to their 401(k).¹⁹

If a high 401(k) return is a sign of high 401(k) investing skill, then we should see persistence in 401(k) portfolio returns over time. We test whether past 401(k) portfolio performance predicts future portfolio performance by regressing an investor's portfolio alpha in year t on her portfolio alpha in year $t - 1$, where $t = 1998, 1999, \text{ and } 2000$. Three-factor alphas are calculated by regressing monthly excess portfolio returns on the excess market return and the Fama and French (1993) size factor and book-to-market factor returns. Four-factor alpha regressions also include Kenneth French's momentum factor (Mom) returns. We cluster our regression standard errors by company \times year \times employee state of residence to account for the fact that asset allocations (and hence alphas) in our sample may not be independently chosen within a company locality. Investors' asset allocations are constrained by the investment options offered in their company's 401(k) plan, so it may be sensible to only compare performance relative to other investors in the same company. Therefore, we also have specifications that include company \times year dummies as explanatory variables.

Table V shows the results of these portfolio performance persistence regressions. We see that, if anything, a good 401(k) portfolio performance this year predicts *poor* performance the

¹⁹ Calvet, Campbell, and Sodini (2007) find that investors stop holding both stocks and bonds after realizing poor mutual fund performance. Because Companies A and E do not offer cash as a 401(k) investment option, one may worry that our performance-chasing results are caused by investors reducing their 401(k) contributions because they simply want to reduce their risky asset share, rather than because they are shying away from 401(k) investing per se. However, our performance-chasing results are robust to restricting the sample to Companies B, C, and D, which offer cash funds in their 401(k).

following year. When we do not control for company \times year effects, 100 basis points of out-performance in year $t - 1$ is associated with 42 basis points of under-performance in year t using three-factor alphas, and 11 basis points of underperformance in year t using four-factor alphas. These effects attenuate when we introduce company \times year dummies in the regression. Three-factor alphas are still negatively serially correlated, while the four-factor alpha exhibits positive serial correlation that is both statistically and economically insignificant. The point estimate indicates that 100 basis points of four-factor alpha are associated with only 2 basis points of four-factor alpha the following year.²⁰

Therefore, there appears to be no empirical support for the hypothesis that returns-chasing and variance-avoidance is driven by rational learning about one's own investing skill.

B. Rebalancing

There is another potential alternative explanation for our finding that 401(k) contribution changes are positively related to portfolio returns: if an investor has significant non-401(k) financial assets, then a positive correlation between 401(k) and non-401(k) asset returns could produce the appearance of return chasing due to rebalancing. For example, suppose all households followed a rule of maintaining a fixed dollar amount in non-401(k) assets (a buffer stock). Then a high 401(k) return would be associated with a high non-401(k) return, which would cause high-return households to increase 401(k) contributions and increase consumption out of non-401(k) assets to bring non-401(k) asset values back down to baseline.

²⁰ The significantly positive constant in the 4-factor alpha regression without company \times year dummies mostly reflects the fact that our sample companies' stocks did extraordinarily well through much of the sample period, and their employees often held significant amounts of those stocks. See Benartzi (2001) for evidence that rank-and-file employees do not have the ability to predict their employer's stock return.

This story, however, is inconsistent with some of our other findings. Such a rebalancing effect should diminish as non-401(k) financial assets get smaller, since the fraction of income required to restore the non-401(k) balance to its steady-state level diminishes for a given percent return. Therefore, the rebalancing story predicts that apparent return chasing would be weakest among the young, who have few financial assets, and strongest among the old. The results presented above in Section III.C, however, showed that the empirical pattern is exactly the opposite: return chasing decreases with age.

Furthermore, most 401(k) households have minimal liquid wealth outside of their 401(k) with which to engage in rebalancing. In the 2001 Survey of Consumer Finances, among 401(k)-holding households earning between \$20,000 and \$70,000 a year—a sample roughly comparable to the one we use in our analysis—the median household has gross non-retirement financial assets equal to only 2.1 months of income, 76% of which is held in checking, savings, or money market accounts.²¹ It is only at the 82nd percentile that households have one year’s income in gross non-retirement financial assets. These figures probably overstate outside asset holdings in our sample because the generosity of our 401(k) plans’ early withdrawal and loan provisions substantially mitigates the need for a precautionary wealth stock outside the 401(k).

Finally, the rebalancing channel cannot explain the robust variance-avoidance we observe among our investors.

V. Conclusion

We have presented evidence that when there is no salient reference purchase price, individual investors chase their own historical returns and shy away from their own historical

²¹ We count CDs, bonds, savings bonds, publicly traded stock, mutual funds, cash value life insurance, other managed accounts, transactions accounts, and miscellaneous assets as non-retirement financial assets.

return variance. Specifically, we find that investors who experience high returns or low variance in their 401(k) portfolio increase their 401(k) contributions more than workers in the same savings plan who experience low returns and/or high variance. This behavior cannot be accounted for by public news about asset returns, investor fixed effects, wealth effects, or time shocks that are correlated with the tendency to hold equities, bonds, or cash. Moreover, we find no evidence that the return chasing and volatility flight are welfare improving, since 401(k) portfolio performance is not persistent. These results are in sharp contrast to the disposition effect, which induces contrarian behavior when there is a salient reference purchase price.

The observed patterns are explained by a naïve reinforcement learning heuristic: assets in which one has *personally* experienced success are expected to be successful in the future. Consistent with reinforcement learning, we also find evidence for the Power Law of Practice: return chasing and volatility avoidance decline with age as a large stock of experience is accumulated, though they remain present throughout the lifecycle.

References

- Agnew, Julie, Pierluigi Balduzzi, and Annika Sundén, 2003. "Portfolio Choice and Trading in a Large 401(k) Plan." *American Economic Review* 93(1), pp. 193 – 215.
- Barber, Brad M., Terrance Odean, and Michal Strahilevitz, 2004. "Once Burned, Twice Shy: Naïve Learning, Counterfactuals, and the Repurchase of Stocks Previously Sold." Mimeo, University of California, Berkeley.
- Barber, Brad M., Terrance Odean, and Lu Zheng, 2000. "The Behavior of Mutual Fund Investors." Mimeo, University of California, Davis.
- Barberis, Nicholas, and Wei Xiong, 2006. "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation." Mimeo, Yale University.
- Benartzi, Shlomo, 2001. "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock." *Journal of Finance* 56, pp. 1747-1764.
- Benartzi, Shlomo, and Richard Thaler, 2001. "Naïve Diversification Strategies in Retirement Savings Plans." *American Economic Review* 91(1), pp.79-98.
- Berk, Jonathan B., and Richard C. Green, 2004. "Mutual Fund Flows and Performance in Rational Markets." *Journal of Political Economy* 1112, pp. 1269-1295.
- Brown, Stephen, and William Goetzmann, 1995. "Performance Persistence." *Journal of Finance* 50, pp. 679-698.
- Calvet, Laurent, John Y. Campbell, and Paolo Sodini, 2007. "Fight or Flight? Portfolio Rebalancing by Individual Investors." Mimeo, Harvard University.
- Carhart, Mark M., 1997. "On Persistence in Mutual Fund Performance." *Journal of Finance* 52, pp. 57-82.
- Charness, Gary, and Dan Levin, 2003. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect." *American Economic Review* 95, pp. 1300-1309.
- Chevalier, Judith, and Glenn Ellison, 1997. "Risk Taking by Mutual Funds as a Response to Incentives." *Journal of Political Economy* 105, pp. 1167-1200.
- Choi, James J., David Laibson, and Brigitte C. Madrian, 2004. "Plan Design and 401(k) Savings Outcomes." *National Tax Journal* 57, pp. 275-298.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2004. "Employees' Investment Decisions About Company Stock." In Olivia S. Mitchell and Stephen P. Utkus, editors, *Pension Design and Structure: New Lessons from Behavioral Finance*. Oxford: Oxford University Press, pp. 121-137.

- Choi, James, David Laibson, and Andrew Metrick, 2002. "How Does the Internet Affect Trading? Evidence from Investor Behavior in 401(k) Plans." *Journal of Financial Economics* 64, pp. 397-421.
- Constantinides, George M., 1984. "Optimal Stock Trading with Personal Taxes: Implications for Prices and the Abnormal January Returns." *Journal of Financial Economics* 13, pp. 65-89.
- Coval, Joshua D., and Tyler Shumway, 2005. "Do Behavioral Biases Affect Prices?" *Journal of Finance* 60, pp. 1-34.
- Dhar, Ravi, and Ning Zhu, 2006. "Up Close and Personal: An Individual Level Analysis of the Disposition Effect." *Management Science* 52, pp. 726-740.
- Duflo, Esther, and Emmanuel Saez, 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118, pp. 815-842.
- Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake, 1996. "The Persistence of Risk-Adjusted Mutual Fund Performance." *Journal of Business* 69, pp. 133-157.
- 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.
- Fama, Eugene F., and Kenneth R. French, 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33, pp. 3-56.
- Feng, Lei, and Mark S. Seasholes, 2005. "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?" *Review of Finance* 9, pp. 305-351.
- Frazzini, Andrea, and Owen Lamont, 2006. "Dumb Money: Mutual Fund Flows and the Cross-Section of Stock Returns." *Journal of Financial Economics*, forthcoming.
- Genesove, David, and Christopher Mayer, 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics* 116, pp. 1233-1260.
- Goetzmann, William, and Roger Ibbotson, 1994. "Do Winners Repeat: Predicting Mutual Fund Performance." *Journal of Portfolio Management* 20, pp. 9-18.
- Goetzmann, William, and Massimo Massa, 2002. "Daily Momentum and Contrarian Behavior of Index Fund Investors," *Journal of Financial and Quantitative Analysis* 37, pp. 375-390.
- Goetzmann, William, and Nadav Peles, 1993. "Cognitive Dissonance and Mutual Fund Investors." *Journal of Financial Research* 20, pp. 145-158.
- Grinblatt, Mark, and Sheridan Titman, 1992. "The Persistence of Mutual Fund Performance." *Journal of Finance* 47, pp. 1977-1984.

- Gruber, Martin J., 1996. "Another Puzzle: The Growth in Actively Managed Mutual Funds." *Journal of Finance* 51, pp. 783-810.
- Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1993. "Hot Hands in Mutual Funds: Short-Run Persistence in Relative Performance, 1974-1988." *Journal of Finance* 48, pp. 93-130.
- 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, 2006. "401(k) plan asset allocation, account balances, and loan activity in 2005." *EBRI Issue Brief* 296.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, 2004. "Social Interaction and Stock-Market Participation." *Journal of Finance* 59, pp. 137-163.
- Huberman, Gur, Sheena Iyengar, and Wei Jiang, 2006. "Defined Contribution Pension Plans: Determinants of Participation and Contribution Rates." Forthcoming in *Journal of Financial Services Research*.
- Huberman, Gur, and Wei Jiang, 2006. "Offering vs. Choice by 401(k) Plan Participants: Equity Exposure and Number of Funds." *Journal of Finance* 61(2), pp. 763-801.
- Huberman, Gur, and Paul Sengmueller, 2004. "Performance and Employer Stock in 401(k) Plans", *Review of Finance* 8, pp. 403-443.
- Ippolito, Richard A., 1992. "Consumer Reaction to Measures of Poor Quality: Evidence from the Mutual Fund Industry." *Journal of Law and Economics* 35, pp. 45-70.
- Ivkovich, Zoran, and Scott Weisbenner, 2007. "Information Diffusion Effects in Individual Investors' Common Stock Purchases: Covet Thy Neighbors' Investment Choices." *Review of Financial Studies*, forthcoming.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance* 48, pp. 65-91.
- Kahneman, Daniel, and Amos Tversky, 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47, pp. 263-291.
- Locke, Peter, and Steven Mann, 2005. "Professional Trader Discipline and Trade Disposition." *Journal of Financial Economics* 76, pp. 401-444.
- Nicolosi, Gina, Liang Peng, and Ning Zhu, 2004. "Do Individual Investors Learn from Their Trading Experience?" Mimeo, University of California, Davis.
- Odean, Terrance, 1998. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance* 53, pp. 1775-1798.

- Patel, Jayendu, Richard J. Zeckhauser, and Darryll Hendricks, 1994. "Investment Flows and Performance: Evidence from Mutual Funds, Cross-Border Investments, and New Issues." In Ryuzo Sato, Richard M. Levich, and Rama V. 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.
- Sapp, Travis, and Ashish Tiwari, 2004. "Does Stock Return Momentum Explain the 'Smart Money' Effect?" *Journal of Finance* 59, pp. 2605-2622.
- Seru, Amit, Tyler Shumway, and Noah Stoffman, 2006. "Learning by Trading." Mimeo, University of Michigan.
- Shapira, Zur, and Itzhak Venezia, 2001. "Patterns of Behavior of Professionally Managed and Independent Investors." *Journal of Banking and Finance* 25, pp. 1573-1587.
- Shefrin, Hersh, and Meir Statman, 1985. "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence." *Journal of Finance* 40, pp. 777-782.
- Sirri, Erik R., and Peter Tufano, 1998. "Costly Search and Mutual Fund Flows." *Journal of Finance* 53, pp. 1589-1622.
- 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.
- Wermers, Russ, 2003. "Is Money Really 'Smart?' New Evidence on the Relation Between Mutual Fund Flows, Manager Behavior, and Performance Persistence." University of Maryland Working Paper.
- Zheng, Lu, 1999. "Is Money Smart? A Study of Mutual Fund Investors' Fund Selection Ability." *Journal of Finance* 54, pp. 901-933.

Table I. 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%	Data unavailable	83%	65%
% married	56%	55%	75%	Data unavailable	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	Allowed	Allowed	Allowed	Allowed
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

Table II. Regression Sample Summary Statistics

This table presents summary statistics on year-over-year change in the contribution rate effective during the last pay cycle of December (ΔC), annual 401(k) percent return (R), annualized 401(k) return variance ($\sigma^2(R)$), and annual 401(k) dollar capital gains normalized by annual income ($CapitalGain/Y$) for the regression sample. The contribution rate change statistics are from year-end 1998 through year-end 2000. The 401(k) return, 401(k) return variance, and normalized capital gains statistics are from year-end 1997 through year-end 2000. Capital gains from 1998 to 2000 are normalized by contemporaneous year income, and capital gains in 1997 are normalized by 1998 income due to the lack of 1997 income data.

	Annual contribution rate change	Annual portfolio return	Annualized portfolio return variance	Annual capital gains/ annual income
Maximum	20%	127.60%	61,935.15	7.87
99 th percentile	9%	81.50%	31,923.33	1.20
90 th percentile	2%	25.61%	11,266.21	0.31
75 th percentile	0%	17.63%	5,223.65	0.13
50 th percentile	0%	9.90%	2,264.68	0.04
25 th percentile	0%	5.48%	698.30	0.01
10 th percentile	0%	-5.47%	1.91	-0.08
1 st percentile	-9%	-26.93%	0.03	-0.56
Minimum	-20%	-82.64%	0.03	-8.49
Mean	0.30%	11.88%	4,418.88	0.09
Std. deviation	2.47%	16.87%	6,306.86	0.30

Table III. Regressions of Contribution Rate Changes on Portfolio Returns and Variance

This table presents coefficients from estimating regression equation (4). The dependent variable is the year-over-year change in the contribution rate effective during the last pay cycle of December. The Δ operator is for year-over-year changes. The subscript i indexes investors, and t indexes years. $R_{i,t}$ is annual 401(k) percent return, $\sigma^2(R_{i,t})$ is annualized 401(k) return variance, $CapitalGain_{i,t}$ is the 401(k) dollar capital gain, $Y_{i,t}$ is annual salary, and $Tenure_{i,t}$ is the number of years since original hire at the end of year t . The last three table rows indicate whether the regression includes company \times year dummies, asset class (equities, bonds, or cash) balances at the prior year-end interacted with year dummies, and the share of the 401(k) in equities or bonds interacted with year dummies. Standard errors, clustered by company \times employee's state of residence in 1998, are in parentheses below the point estimates.

$\Delta R_{i,t}$	0.0078** (0.0012)	0.0078** (0.0016)	0.0083** (0.0016)	0.0071** (0.0022)
$\Delta R_{i,t-1}$	0.0009 (0.0015)	-0.0002 (0.0025)	-0.0010 (0.0027)	-0.0010 (0.0027)
$\Delta\sigma^2(R_{i,t})/1000$	-0.0288** (0.0061)	-0.0288** (0.0061)	-0.0301** (0.0068)	-0.0259** (0.0065)
$\Delta\sigma^2(R_{i,t-1})/1000$	-0.0269** (0.0100)	-0.0267** (0.0101)	-0.0311* (0.0125)	-0.0275* (0.0118)
$\Delta(CapitalGain_{i,t}/Y_{i,t})$		0.0207 (0.0792)	-0.2442* (0.0992)	-0.1848 (0.0991)
$\Delta(CapitalGain_{i,t-1}/Y_{i,t})$		0.1837 (0.2165)	0.5439 (0.2959)	0.5731 (0.2944)
$\Delta\text{Log}(Tenure_{i,t})$	-1.1662 (0.9470)	-1.2709 (0.9964)	-1.0890 (0.9755)	-1.2208 (0.9942)
Company \times Year dummies	Yes	Yes	Yes	Yes
Balance \times Year controls	No	No	Yes	Yes
Share \times Year controls	No	No	No	Yes

* Significant at the 5% level. ** Significant at the 1% level.

Table IV. Regressions of Contribution Rate Changes on Portfolio Returns and Variance Interacted With Age and Income

This table presents coefficients from estimating a variant of the regression equation (4). The dependent variable is the year-over-year change in the contribution rate effective during the last pay cycle of December. The Δ operator is for year-over-year changes. The subscript i indexes investors, and t indexes years. $R_{i,t}$ is annual 401(k) percent return, $Age_{i,1998}$ is age at year-end 1998, $Y_{i,1998}$ is 1998 salary, and $\sigma^2(R_{i,t})$ is annualized 401(k) return variance. All regressions control for contemporaneous and lagged 401(k) dollar capital gains normalized by annual income ($CapitalGain/Y$), tenure at company, company \times year dummies, asset class (equities, bonds, or cash) balances at the prior year-end interacted with year dummies, and the share of the 401(k) in equities or bonds interacted with year dummies. Standard errors, clustered by company \times employee's state of residence in 1998, are in parentheses below the point estimates.

$\Delta R_{i,t}$	0.0059** (0.0022)	0.0071** (0.0023)	0.0059** (0.0023)
$(\Delta R_{i,t}) \times Age_{i,1998}/10$	-0.0024 (0.0014)		-0.0023 (0.0014)
$(\Delta R_{i,t}) \times \log(Y_{i,1998})$		-0.0007 (0.0027)	-0.0010 (0.0028)
$\Delta R_{i,t-1}$	-0.0016 (0.0025)	-0.0020 (0.0030)	-0.0026 (0.0027)
$(\Delta R_{i,t-1}) \times Age_{i,1998}/10$	-0.0005 (0.0040)		-0.0003 (0.0041)
$(\Delta R_{i,t-1}) \times \log(Y_{i,1998})$		-0.0117 (0.0072)	-0.0122 (0.0077)
$\Delta \sigma^2(R_{i,t})/1000$	-0.0224** (0.0055)	-0.0255** (0.0066)	-0.0221** (0.0057)
$(\Delta \sigma^2(R_{i,t})/1000) \times Age_{i,1998}/10$	0.0196** (0.0040)		0.0200** (0.0042)
$(\Delta \sigma^2(R_{i,t})/1000) \times \log(Y_{i,1998})$		0.0005 (0.0100)	-0.0020 (0.0111)
$\Delta \sigma^2(R_{i,t-1})/1000$	-0.0253* (0.0111)	-0.0271* (0.0117)	-0.0248* (0.0111)
$(\Delta \sigma^2(R_{i,t-1})/1000) \times Age_{i,1998}/10$	0.0229** (0.0059)		0.0230** (0.0059)
$(\Delta \sigma^2(R_{i,t-1})/1000) \times \log(Y_{i,1998})$		0.0175 (0.0132)	0.0150 (0.0138)
<i>CapitalGain/Y</i> controls	Yes	Yes	Yes
Tenure controls	Yes	Yes	Yes
Company \times Year dummies	Yes	Yes	Yes
Balance \times Year controls	Yes	Yes	Yes
Share \times Year controls	Yes	Yes	Yes

* Significant at the 5% level. ** Significant at the 1% level.

Table V. 401(k) Return Performance Persistence

This table shows the results of regressing year t 401(k) portfolio alpha on year $t - 1$ 401(k) portfolio alpha, where t goes from 1998 to 2000. The 3-factor alpha controls for the market return, size effect, and book-to-market effect. The 4-factor alpha also controls for stock price momentum. The last row indicates whether the regression includes company \times year dummies. Standard errors, clustered by company \times year \times employee's state of residence in year t , are in parentheses below the point estimates.

	3-factor alpha	3-factor alpha	4-factor alpha	4-factor alpha
α_{t-1}	-0.4226** (0.0678)	-0.1722** (0.0520)	-0.1070 (0.0784)	0.0213 (0.0614)
Constant	0.1495 (0.1104)	--	0.3849** (0.0994)	--
Company \times Year dummies	No	Yes	No	Yes