Do Fluctuations in Wealth Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation

Markus Brunnermeier*
Princeton University and CEPR

Stefan Nagel+
Stanford University

PRELIMINARY

November 2004

We thank Frank de Jong, Jonathan Parker and workshop participants at the London School of Economics and Stanford for useful comments.
* Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, email: markus@princeton.edu; http://www.princeton.edu/~markus
+ Stanford University, Graduate School of Business, 518 Memorial Way, Stanford, CA 94305, email: Nagel_Stefan@gsb.stanford.edu; http://faculty-gsb.stanford.edu/nagel/

ABSTRACT: One explanation for the apparent "excess volatility" and predictability of stock returns is that fluctuations in wealth generate counter-cyclical variation in risk aversion, as, for example, in models with additive habit formation preferences. While these representative-agent models have some success in matching moments of aggregate variables, it is not clear yet what their micro-foundations are. In an effort to provide evidence on this issue, we analyze two decades of micro data from the PSID and CEX surveys to estimate how a typical household's willingness to bear stock market risk responds to wealth shocks. Our results show that there is no positive relationship between changes in household wealth and the share of financial wealth allocated to stocks. If anything, the relationship is slightly negative. Instead, we find that the dominant influence on individuals' asset allocation is inertia: Following a capital gain or loss or in- and outflows of financial wealth, households do very little rebalancing. But even controlling for this inertia, there is no economically significant wealth effect on stock holdings. Overall, our results suggest that the apparent negative relationship between wealth changes and risk aversion at the aggregate level is not driven by a similar relationship at the household level. Moreover, capital gains or losses on stocks appear to play a special role in explaining household asset allocation.

1. Introduction

The equity volatility puzzle—that is, the predictability and “excess volatility” of stock market returns—is one of the major unresolved issues in asset pricing (Campbell 2000). Expected stock returns vary counter-cyclically, thereby generating much of the volatility in realized returns. Recent work in macroeconomics and finance has attempted to resolve this puzzle by allowing the representative agent’s risk aversion to respond negatively to wealth shocks. To take one prominent example, in Campbell and Cochrane (1999), with additive habit formation preferences, a positive wealth shock decreases the local curvature of the utility function, and thus, the relative risk aversion of the representative agent. As a result, there is counter-cyclical variation in the conditional equity risk premium. Similar wealth effects can arise when agents worry about relative social status as in Bakshi and Chen (1996), or when they have the generalized disappointment aversion preferences of Routledge and Zin (2003).

While these representative-agent models are relatively successful in matching aggregate time-series data, little is known yet about their micro foundations. An appealing way to justify these preferences with time-varying risk aversion at the aggregate level would be to have individuals with similar preferences at the micro-level. A direct of implication of this theory would be that individuals’ relative risk aversion should decrease in response to positive wealth shocks. To the extent that the wealth shocks are idiosyncratic, such shifts in risk aversion in turn should influence their allocation of financial wealth between risky and riskless assets. Whether wealth shocks do indeed affect household asset allocation in this way has not been investigated yet in the existing literature.

Our objective in this paper is to explore empirically whether habit formation—or alternative forms of time-varying risk aversion based on wealth effects—at the individual level can provide a micro foundation for representative-agent habit models. To this end, we estimate to what extent changes in wealth alter households’ willingness to bear stock market risk. If relative risk aversion responded negatively to wealth shocks, there should be a positive relationship between shocks to wealth and changes in the share of financial wealth allocated to risky assets. Moreover, based on the Campbell-Cochrane calibrations, we show that, to be consistent with the large variation in conditional expected returns at the aggregate level, this effect would have to be rather large and easily detectable in micro data.

We use household level data from the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX), covering the time period from the mid-1980s to 2001. We construct measures of total wealth, including home equity, and the share of risky assets (stocks and mutual funds)[1] in their portfolios of financial wealth. The PSID is longitudinal, and we have observations on each household's asset allocation and wealth at intervals of several years. The CEX is a short panel, and we can follow each household over four quarters. In all our regressions, we condition on past stock market participation of the household, because the decision to participate in the stock market should be seen as separate from the question of how to allocate between risky and riskless assets, conditional on past participation (see, e.g, Vissing-Jorgensen 2002).

Our main result is easily summarized. We find, unambiguously, both in the PSID and CEX, that there is no positive relationship between wealth shocks and changes in the fraction of financial wealth allocated to risky assets. This is not the result of low statistical power—our coefficients are quite precisely estimated. If anything, the effect is slightly negative. We control for predictable movements in wealth and asset allocations, in particular life-cycle effects, with age group, wealth group, and year effects, as well as for exposure to housing wealth, equity in private business, and a variety of other aspects. Also, using household consumption growth in place of wealth changes delivers the same result. The result also holds up when we control for slow adjustment in portfolio allocations. This is important, because one might conjecture that increases in wealth might tend to materialize first in relatively liquid form, e.g., on checking or savings

[1] In the CEX data, the set of risky assets also includes government bonds (but not savings bonds). However, only few individuals invest in government bonds at all, so, in practice, this does not make a substantial difference.

accounts. With infrequent or delayed adjustment, this could result a mechanical negative association of changes in wealth and the proportion of financial wealth allocated to stocks. After controlling for such inertia, our estimate of the wealth effect is still basically equal to zero. This suggests that a typical household's relative risk aversion does not change in response to wealth shocks.

While wealth shocks do not seem to affect asset allocation at all, inertia, in contrast, has a dramatic impact. It seems that, to a large extent, households do not rebalance following capital gains or losses on their risky asset portfolios. Even five years into the future, past capital gains still have a substantial effect on asset allocations. Underreporting of trades in the survcy data may also contribute to this effect, but it is unlikely to account for all of it. In particular, the strong inertia we find parallels similar findings in 401(k) retirement account data by Samuelson and Zeckhauser (1988), Ameriks and Zeldes (2001), Agnew, Balduzzi, and Sunden (2000), and Huberman and Sengmueller (2004). It suggests that household portfolio allocations depend on the path of asset prices experienced by the individual household, and are therefore quite idiosyncratic.

Overall, our results cast doubt on the hypothesis that variation in relative risk aversion driven by wealth shocks can provide a micro foundation for aggregate time-varying risk aversion models. This effect does not appear to exist at the micro level. Our findings using a first-differences model—which is the appropriate framework for the question we focus on—are also consistent with earlier evidence that, conditional on stock market participation, the cross-sectional relationship between the *level* of the risky asset share and the *level* of wealth is essentially flat (Heaton and Lucas 2000; Guiso, Haliassos, and Jappelli 2003).

Of course, this is does not rule out that there could exist a micro foundation of some other form. For example, from Constantinides and Duffie (1996) and Mankiw (1986) we know that one can generate *any* marginal utility process for the representative agent with heterogeneous power utility individuals, incomplete markets, and a judiciously chosen process for the cross-sectional variance of idiosyncratic income. Alternatively, what looks like time-varying risk aversion at the macro level could perhaps be due to distorted beliefs at the micro level (e.g., Campbell and Kyle 1993; Cecchetti, Lam, and Mark 2000; Campbell and Vuolteenaho 2004). Alas, to predict (in a structural sense) when and how expected stock market returns change, we need to understand what drives individuals' willingness to bear stock market risk at the micro level.

Our findings on inertia suggest that capital gains and losses play a special role in determining households' allocations to risky assets. For ease of reference, we labeled this effect as "inertia". It could be due to true inertia, driven by a combination of transaction costs, cognitive costs, and limited attention, but it could also arise because the experience of capital gains and losses has a special impact on households' preferences with respect to stock market risk (as in models with narrow framing of risks), or on beliefs about future returns (trend-chasing). We discuss some of these potential explanations at the end.

Our analysis is related to other research that connects asset pricing theory and micro data. The focus of existing work is mostly on the unconditional equity premium, while our focus is on the time-variation in risk aversion. Mankiw and Zeldes (1991), Parker (2001), Brav, Constantinides, and Geczy (2002), Vissing-Jorgensen and Attanasio (2002) use the average consumption of agents with non-zero stock holdings to test consumption-based pricing models. Heaton and Lucas (2000) investigate the effect of entrepreneurial income risk. Barsky et al. (1997) examine how measures of risk tolerance obtained form survey questions in the Health and Retirement Study relate to wealth and observed portfolio choices, and Dynan (2000) tests the implications of (internal) habit formation for household consumption. The only paper looking at micro-level asset allocation in the context of habit formation that we are aware of is Lupton (2003), who finds a negative relationship between past consumption levels and current risky asset holdings (in $), which he interprets as being consistent with habit formation. Our approach leads us to different conclusions for two reasons. Unlike his levels specification, our first-differences model allows us to control for unobserved heterogeneity, which is crucial in micro data. Looking at wealth shocks rather than levels also makes economic sense the habit formation story because we do not need to specify the law for the evolution of habit to test for wealth effects.

The paper is organized as follows. Section 2 presents a simple partial equilibrium exercise to assess the magnitude of the wealth effects that we should expect to find if individuals have habit formation preferences consistent with the calibration in Campbell and Cochrane (1999). Section 3 describes the data and presents some summary statistics. In Section 4, we show how we estimate wealth effects on asset allocation using either wealth or consumption data, and we present our estimation results. In Section 5 we use regressions that control for inertia effects, and we distentangle inertia with respect to capital gains and in-/outflows of financial wealth. Section 6 discusses the implications of our findings for asset pricing theory.

2. Theoretical Predictions and Econometric Framework

2.1 Quantitative implications of additive external habits

Before getting into our empirical analysis, we want to assess the order of magnitude of the wealth effects predicted by models that match the time-variation in expected returns found in historical data. This will tell us whether we can at all hope to detect the effect in micro data, or whether it is likely to be swamped by estimation error. For concreteness, we focus on the Campbell and Cochrane (1999) model, using parameters from their calibrations. However, the point should be more general, because, to achieve the same variation in conditional expected returns, other models, calibrated to match the same aggregate asset return and consumption processes, would have to generate movements in utility curvature of the same order of magnitude from the same wealth shocks.

Consider an agent who maximizes expected utility

$$\mathrm{E}\sum_{t=0}^{\infty}\delta^t\,\frac{\left(\mathrm{C}_t-\mathrm{X}_t\right)^{1-\gamma}-1}{1-\gamma}\,, \qquad (1)$$

where X_t is the habit, and δ is the subjective time discount factor. Define the surplus-consumption ratio as $S_t \equiv (C_t\text{-}X_t)/C_t$. To allow for heterogeneous wealth levels, we can think of X_t as being determined by the average consumption of a reference group. Within this reference group, agents are identical. But across reference groups, there can be differences in levels of C_t and expected consumption growth, and reference groups may get idiosyncratic wealth shocks. These idiosyncratic shocks will make some groups want to hold more of their wealth in risky assets, while others will want to hold less in risky assets. As in Campbell and Cochrane, habit is assumed to be external, and thus, an individual agent does not consider how current consumption affects future habits. Furthermore, for each reference group, the surplus-consumption ratio follows the same nonlinear law as in Campbell and Cochrane, with the same parameters.

Now consider an agent who is initially at the steady-state value $S_t = \overline{S}$ and who subsequently experiences consumption growth of $\Delta c_{t+1} \equiv c_{t+1}\text{-}c_t = g+\varepsilon_{t+1}$, where g is expected consumption growth and ε_{t+1} is an idiosyncratic shock. In this case, the surplus-consumption process in Campbell-Cochrane implies that Δc_{t+1} and $\Delta s_{t+1} \equiv \log(S_{t+1}) - \log(S_t)$ are related as follows:

$$\Delta s_{t+1} = \frac{1-\overline{S}}{\overline{S}} \varepsilon_{t+1}. \tag{2}$$

Furthermore, Campbell and Cochrane (1999) show that local utility curvature (relative risk aversion) is inversely related to S_t:

$$\eta_t = \frac{\gamma}{S_t}. \tag{3}$$

The agent's problem is to allocate wealth between these a risky asset with log return r_{mt}, and a risk-free asset with log return r_{ft}. To get a simple closed-form approximation of the solution, we assume joint conditional log-normality and homoskedasticity of asset returns and consumption, as in Hansen and Singleton (1983), Campbell (1993). Furthermore, we ignore hedging demands for a moment and focus on the portfolio choice of a myopic investor (or, equivalently, on the case of constant investment opportunities), which yields the standard expression for the optimal share of the risky asset in the agents' portfolio of

$$Q_t = \frac{E_t[r_{mt+1} - r_{ft+1}] + \sigma_m^2/2}{\eta_t \sigma_{mc}}, \tag{4}$$

where $\sigma_{mc} \equiv Cov[r_{mt}, \Delta c_t]$ and $\sigma_m^2 \equiv Var[r_{mt}]$. Substituting (3) into (4), and defining $\Delta\rho_t \equiv \log(Q_t/Q_{t-1})$, it is easy to see that $\Delta\rho_{t+1} = \log(S_{t+1}/S_t)$, and hence, by equation (2), if the agent is initially in steady-state at time t, and faced with an unexpected consumption shock ε_{t+1}, the optimal risky asset share changes as follows:

$$\Delta\rho_{t+1} = \frac{1-\bar{S}}{\bar{S}}\varepsilon_{t+1}. \tag{5}$$

The intuition is simple. When C_t is above, but close to habit in the steady state, and hence $\bar{S}$ is low, a given change in consumption leads to a large change in local curvature, and hence a large change in the optimal portfolio share of risky assets.

Eq. (5) now provides us with an idea of the magnitudes involved. Let $\bar{S} = 0.057$, as in the Campbell-Cochrane calibrations. In this case, a small unexpected idiosyncratic shock to consumption of 1% would imply a sizeable increase in the risky asset share by 16.5% (e.g., from 50% to 58.25%). If the agent consumes a constant fraction of wealth, then the consumption shock ε_t can be replaced by a shock to log wealth, with similar quantitative implications.[2] Since growth in wealth is easier to measure in household-level data than consumption growth, the wealth-version of Eq. (5) will be the main focus of our empirical tests.

If the agent is not at the steady state initially, then Eq. (5) to some extent overstates the wealth-sensitivity of the risky asset share. However, we ran simulations where different reference groups get random consumption shocks over time, and hence, at a given point in time some have S_t

[2] In fact, in this habit formation model, the consumption-wealth elasticity is above one. This means that the sensitivity of $\Delta\rho_t$ to wealth shocks would be higher than to consumption shocks.

below and some above $\overline{S}$. This reduces the wealth-sensitivity only marginally. With $\sigma_c = 1.5\%$, an OLS regression of $\Delta\rho_t$ on Δc_t in the simulated sample still yields a slope of about 14 (instead of 16.5). Also, adding measurement noise to Δc_t with the same variance as Δc_t attenuates the coefficient to about 7. Incorporating hedging demands would also alter the results somewhat. Hedging demands would typically dampen the response of the optimal risk asset share to a change in local curvature (as long as $\eta_t > 1$). Yet, all we want to take away from this simple exercise is that the order of magnitude of the wealth effect should be large—the coefficient in a regression of $\Delta\rho_t$ on Δc_t should at least be well above one, even with noisy measurement of consumption growth. This conclusion is unlikely to be changed by incorporation of hedging demands.

2.2 Specification and identification

Our analysis above implies that additive external habit formation models—or models with similar wealth effects—predict that the share of risk asset holding should change with a shock to wealth. In levels, these theories imply the following infinite distributed lag structure for the log risky asset share:

$$\rho_{it} = \delta_0 + \delta_1\,\varepsilon_{it} + \delta_2\,\varepsilon_{it-1} + \delta_3\,\varepsilon_{it-2} + \ldots + \lambda_t + c_i + z_{it} + \xi_{it}\,, \qquad (6)$$

where $\delta_1 > \delta_2 > \delta_3 \ldots$. Thus, the wealth shocks ε_{it} cause changes in ρ_{it}, but their effect dies out over time, as relative risk aversion reverts back to its steady state value. In addition, λ_t captures effects that are common across individuals, for example, shifts in the preferences of the individuals in our data sets relative to other market participants, c_i captures unobserved time-constant heterogeneity in preferences (e.g. in γ) across individuals, and z_{it} includes time-varying exogenous factors such as life-cycle effects or changes in family composition whose effects are not explicitly modeled here, but which will be important controls. Taking first differences removes c_i, yielding

$$\Delta\rho_{it} = \alpha_0 + \delta_1\,\varepsilon_{it} + \alpha_2\,x_{it} + u_{it}\,,\ \text{ where } u_{it} = \Delta\xi_{it} - \delta_1\,\varepsilon_{it-1} + \delta_2\,\Delta\varepsilon_{it-1} + \delta_3\,\Delta\varepsilon_{it-2} + \ldots\,. \qquad (7)$$

Note that $\Delta\lambda_t$ and Δz_{it} are collapsed into x_{it}, a vector of household characteristics and time dummy variables. This first difference equation is a natural specification to test the implications of habit formation preferences, as it maps directly into Equation (5).

To analyze identification, we need to specify the dynamics of household wealth. We assume that wealth follows a random walk, after controlling for life-cycle effects and other predictable or exogenous wealth movements captured by x_{it}. Hence, in first differences,

$$\Delta w_{it} = \varepsilon_{it} + \theta\, x_{it}. \tag{8}$$

A common conditional expected growth rate g_t is allowed for, as any effects common across individuals would be picked up by the time period dummy variables in x_{it}. Solving Eq. (8) for ε_{it}, substituting into Eq. (7), and denoting $\beta_0 = \alpha_0$, $\beta_1 = \delta_1$, and $\beta_2 = \alpha_2 - \theta$ then leads to

$$\Delta\rho_{it} = \beta_0 + \beta_1\, \Delta w_{it} + \beta_2\, x_{it} + u_{it}\ , \tag{9}$$

which is our baseline regression model. Under the random walk assumption, $\Delta\rho_{it}$ cannot feedback into future Δw_{is}, $s > t$. Moreover, the lags of ε_{it} appearing in u_{it} are uncorrelated with Δw_{it}. Thus, both Δw_{it} and x_{it} are strictly exogenous, i.e., $E[u_{it}|\Delta w_{i1}, \Delta w_{i2}, \ldots, \Delta w_{iT}, x_{i1}, x_{12}, \ldots, x_{iT}] = 0$.[3] In the absence of measurement error, our random walk assumption (8) therefore implies that we can consistently estimate Eq. (9) with OLS. There is reason to believe that after accounting for life-cycle effects, a random walk for wealth is a reasonably good approximation. Note also that if transitory movements in wealth are driven by mean-reversion in aggregate asset returns, our use of time dummies in x_{it} would eliminate these effects. For these reasons, we estimate our baseline regressions with OLS.

An alternative would be to assume an AR(1) model for wealth levels. However, the fact $\Delta\rho_t$ depends on an infinite distributed lag of past ε_{it}, subsumed into the disturbance in Eq. (9), means that we cannot employ the usual approach of using lags of wealth as instruments to get a

[3] In some of our specifications, though, there may be controls that may not be strictly exogenous. It is possible, for example, that changes in risky asset allocations might anticipate a change in the number of children in the family, which is one of our controls. Since its coefficient is close to zero, this is unlikely to have any impact.

consistent estimator. Alternatively, one could allow for transitory variation in wealth by modeling the permanent wealth component as a random walk

$$w_{it}^{P} = w_{it-1}^{P} + \varepsilon_{it} , \tag{10}$$

and the transitory component as a moving average (MA) process. In the simplest case, we could assume that it is serially uncorrelated noise (which could also contain measurement error),

$$w_{it}^{T} = \nu_{it}. \tag{11}$$

In this case, the disturbance in Eq. (9) would include ν_{it}, and hence be correlated with $\Delta w_{it} = \Delta w_{it}^{P} + \nu_{it}$. As a result, the OLS estimator would be inconsistent. Also, lags of w_{it} are still invalid as instruments. However, if we have a second measurement of the permanent wealth shock ε_{it}, consistent estimation may be possible. One candidate measure is consumption growth. Theoretically, it should be related almost one-to-one to changes in permanent wealth. It is also the variable that is directly modeled in consumption-based asset pricing models. But in household-level data, measurement error is large. Suppose that

$$\Delta c_{it} = \varepsilon_{it} + \psi_{it} , \tag{12}$$

where ψ_{it} is serially uncorrelated measurement error (more generally, it could also have an MA structure) which is also uncorrelated with ν_{it}. More precisely, what we need in the case of MA(q) transitory wealth is that $\mathrm{Cov}(\psi_{it}, \nu_{is}) = 0$ for $t-q \le s \le t+q$, which is not an unreasonable assumption. Then, we can put Δc_{it} in place of Δw_{it} in Equation (9) and use Δw_{it} as an instrument for Δc_{it} (or vice versa) to consistently estimate β_1. Of course, given the large measurement error in consumption growth, it may be that the empirical relationship between Δc_{it} and Δw_{it} is too weak to yield reliable estimates in the second stage due to the weak instruments problem analyzed by Staiger and Stock (1997). Trading off robustness against bias, we stick to OLS in most of our wealth regression specifications, but we also explore this IV strategy.

3. Data and Summary Statistics

We draw on two different data sources: the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure survey (CEX). They complement each other in various ways. The PSID data has better quality measures of wealth and asset allocation, but it is weak on capital gains. The CEX data on wealth, in contrast, is not as detailed as in the PSID, and provides a less satisfactory definition of the set of risky assets, but it is sampled at higher frequency and provides more reliable capital gains data.

3.1 Panel Study of Income Dynamics

The PSID, obtained from the University of Michigan, is a longitudinal study that tracks family units and their offspring over time. From 1968 to 1996, the PSID interviewed subjects every year, but it switched to biennial data collection in 1997. Currently, the last available wave is 2001. Annual sample sizes range from 5,000 to 7,000. The data on wealth that we use in this study was collected only in the years 1984, 1989, 1994, 1999, and 2001. We use the core family files from these years to construct all of our variables.

Our definitions are as follows. We define total wealth, W_{it}, as the sum of home equity (value of the home minus remaining mortgage principal), equity in other real estate, equity in a farm or business, equity in vehicles plus financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money market funds, certificates of deposits, savings bonds, or treasury bills), bonds and life insurance (bonds, bond funds, cash value in a life insurance, valuable collection for investment purposes, rights in a trust or estate), stocks (shares of stock in publicly held corporations, mutual funds, investment trusts),[4] minus other debts (such as credit cards, student loans, medical or legal bills, or loans from relatives). From these data we compute

[4]In PSID, subjects are asked to report securities holdings net of amounts owed on the position. In the CEX data below, this is not explicit in the questionnaire, but it is possible that some subjects might interpret it this way.

the share of stocks in the portfolio of financial assets, and we use ρ_{it} to denote its natural logarithm. As the denominator we use financial wealth and add back "other debts".

Before 1999, subjects have been asked explicitly to include assets held in individual retirement accounts (IRA) when reporting their financial asset holdings. Starting in 1999, they have been asked to exclude assets in employer-based pensions and IRAs. Instead, there is a separate question on the value of IRA assets and their allocation to different asset classes. Based on the answer to the latter question, we allocate the IRA assets to stocks and bonds. If subjects state "mostly stocks" we allocate 100% of the IRA value to stocks, if the answer is "split" we allocate 50% to stocks and 50% to bonds, if it says "mostly interest bearing" we put 100% to bonds.

To be included in our sample, we require that the marital status of the family unit head remained unchanged and that no assets have been moved out or in as a consequence of a family member moving out or into the family unit. We further require lagged total wealth and lagged stock holdings greater than zero in the t-1 wave (with waves we always mean those with wealth data only). This means that we are looking at changes in portfolio allocation conditional on past participation. We also delete observations for which one of the wealth components has been topcoded[5] or where the data is only given in brackets. To make magnitudes comparable over time, we deflate wealth by the consumer price index (CPI) into December 2001 dollars.

In years when the wealth questions were administered, the PSID asked subjects to report on the amount of stocks bought and/or sold during the time since the previous wealth survey (i.e., in the 1989 wave for the time from 1984 to 1989; in the 2001 wave for the time from 1999 to 2001). This information allows us to decompose the change in the value of stock holdings into an active investment/disinvestment component and a capital gains component, with its log denoted r_{it}. For IRA assets in 1999 and 2001, we only have a combined active investment figure for all IRA

[5] To assure confidentiality in surveys (or, in some cases, to limit record length), wealth or income observations larger than some benchmark level are commonly replaced by the average of all observations above the benchmark level.

assets. As an approximation, we assume that new IRA funds are allocated pro rata among the prior holdings. Unfortunately, there is reason to expect that this active investment information is noisy. When households are asked to recall their investments over five years, it is to be expected that the information will not be very accurate. In particular, Vissing-Jorgenson (2002) suggests that there is "recall bias" in the PSID: The number of households with zero trades seems too large. For example, some households move from zero stock holdings to becoming a stockholder, but without reporting any investment. To mitigate part of this problem, we include zero-investment observations only if subjects explicitly stated that their buys and sells are equal, or both zero.

Although we do not need capital gains information for our main tests on wealth effects, but only for some further exploration of our results, we make some attempts to mitigate these measurement problems. First, we also employ a second data set, described below, which is administered annually and should therefore suffer less from recall bias. Second, in some of our tests, look at a subsample of those households that report non-zero trades. Third, we calculate percentage capital gains of stocks from t-1 to t and we trim our sample at +/- 2 cross-sectional standard deviations of log capital gains relative to the log market return over the same period.[6] In order to compute capital gains, we need to make an assumption regarding the timing of investment. The reported investment could either have occurred early or late in the measurement period. We assume that half of it has been made at the beginning of the period, and half of it at the end.

3.2 Consumer Expenditure Survey

Our second source of asset allocation data for households is the Consumer Expenditure Survey (CEX) of the Bureau of Labor Statistics, obtained through the Inter-University Consortium for Political and Social Research at the University of Michigan. The CEX has been carried out since 1980, and it is a short panel based on a stratified random sample of the U.S. Population. As

[6] Errors in capital gains are really worrisome only for the inertia effects we analyze below – not for wealth effect tests. But even there doing our estimation with our without the filter has little impact on the regression results.

part of this survey, households are interviewed every three months over a five-quarter period. Then they leave the sample. Each month some new households enter the sample replace those leaving the sample. Most important for our purposes, income and demographic information is collected in the second and fifth interview. Information on financial asset holdings and changes in these holdings over the preceding twelve months is collected in the fifth interview. We extract most of the variables from the family files. Only the data on housing and credit are taken from the detailed expenditure files.

We define total wealth (W_{it}) as the sum of home equity (sum of property values minus sum of outstanding mortgage balances) and financial wealth. Financial wealth is defined as the sum of balances in checking accounts, savings accounts, savings bonds, money owed to the household, and securities (stocks, mutual funds, private bonds, government bonds), minus other debt. Unfortunately, before 1988, there is no information on the level of mortgage balances, only on newly taken up mortgages and repayments. For this reason we use the 1988 to 2001 data only. Also, while we have changes in financial assets over 12 months, for some of the other wealth components (home equity and "other debt"), we can only compute the change over the time from the second to the fifth interview, which is only a nine-month period and hence not perfectly aligned with the financial wealth data. Unlike the PSID, the CEX also does not contain information on holdings in retirement accounts (there is only information on flows), but the questions about security holdings also do not explicitly ask subjects to exclude retirement assets. Hence, it is unclear whether some respondents might include them. For these reasons, the wealth data in the CEX is likely to be of lower quality than in the PSID.

Moreover, the portfolio shares we can compute with the CEX data are different from those in the PSID. The CEX lumps together stock and mutual fund investment with government bonds (except savings bonds) in one position, which we refer to as "securities". In our analysis of CEX data, we will use this as the holdings of risky assets. The inclusion of bonds in this position is certainly not ideal. Yet, Ameriks and Zeldes (2001) report (based on the 1998 Survey of Consumer

Finances) that only 3.2% of individuals hold bonds directly, compared with 20.8% for savings bonds. This suggests that viewing the “securities” position as risky assets is not likely to imply a large error. From these data we compute the share of securities in the portfolio of financial assets. As the denominator we use financial wealth to which we add back “other debts”.

In the fifth interview, subjects in the CEX are asked about the amount of securities purchased and sold over the preceding 12 months. This information allows us to decompose the change in the value of stock holdings into an active investment/disinvestment component and a capital gains component. Since the recall period is much shorter in the CEX (12 months) than in the PSID (60 months), the CEX capital gains data should be more accurate than the PSID data and should allow for a useful check on the PSID results.

Similar to our treatment of the PSID, we also employ some filters to screen out errors. We require that from the second to fifth interview, the marital status of the respondent and the size of the family remained the same, and there is only one consumer unit (family) in the household. As in the PSID, we further require lagged total wealth and lagged stock holdings greater than zero in the second interview. We also delete observations for which one of the wealth components has been topcoded. In the CEX, this is happening more frequently than in the PSID. Finally, to screen out likely data errors in the capital gains measure, we calculate percentage capital gains of stocks from t-1 to t and we trim our sample at +/- 2 cross-sectional standard deviations of log capital gains relative to the log market return over the same period. The assumption about investment timing when computing capital gains is the same as above in the PSID case.

3.3 Summary statistics

Table 1 presents summary statistics for the PSID sample. The first set of columns shows statistics for the set of all family units that have all the required data on wealth components and also lagged wealth greater than zero. The second set of columns looks at the subset of family units that participated in the stock market in the previous wave of the survey and have a valid observation for the capital gains variable (i.e., that satisfy our filter). Looking at the number of observations, one can see that conditioning on past stock market participation reduces the sample substantially. As is well known from previous studies (e.g., Vissing-Jorgensen 2002), only a relatively small fraction of households hold stock, but with an upward trend during the 1980's and 1990's. The number of stockholders here is somewhat understated, though, because the capital gains filter that is imposed on the participant sample. The percentage invested in risky assets (%risky) has gone up over time for both the full and the participant sample. The two right-most columns show cross-sectional statistics for the two main variables in our empirical tests, Δw_{it}, the change in log real wealth, and $\Delta\rho_{it}$, the change in the log portfolio share of risky assets. They show that both wealth and risky asset holdings for the typical household have gone up over the sample period, but also that there is large cross-sectional variation—as shown by the cross-sectional standard deviation in parentheses—in wealth and risky asset share changes. The small growth in wealth from 1989 to 1994 is consistent with the fact that this period contains a recession.

The bottom panel presents the same statistics, but with the sample broken down into quartiles of lagged wealth instead of by year. For the full sample, it seems like there is a strong upward sloping relationship between wealth and the portfolio share of risky assets. However, a comparison with the stock market participant sample shows that this effect is largely due to a strong positive relationship participation and wealth level. Conditional on participation, the relationship between risky asset holdings and wealth is mostly flat. Similar results have been found with data from the Survey of Consumer Finances by Heaton and Lucas (2000), and, with international data, by Guiso, Haliassos, and Jappelli (2003). This finding of a flat relationship between wealth and risky asset holdings in levels already suggests some doubt about the existence of wealth effects on asset allocation. After all, if the typical investor responded to wealth shocks with changes in risk aversion and asset allocation, one might expect that this should also lead to at least some dependence of the portfolio share of risky assets on the level of wealth. Of course, it is important to control for heterogeneity of households along many other dimensions, which is why we analyze this question in first differences instead of levels.

Importantly, the wealth quartile statistics also show some tendency of real wealth to mean-revert: Δw_{it} for the lowest wealth quartile is positive and large, while much smaller for the highest lagged wealth quartile. Two likely explanations for this phenomenon are life-cycle effects and measurement error.[7] Our regression specifications will control for such predictable movement in wealth by including dummies and interaction terms for lagged wealth levels and age.

Table 2 reports similar statistics for the CEX. The number of observations shows a strong drop in 1996 because of changes in the survey that make it impossible to track households across the year-end of 1995. The stock market participation rates and risky asset holdings are similar to those from the PSID. This also shows that the difference in the definition of risky asset holdings (for the CEX it includes treasury bonds) does not make much of a difference. For example, in 1989 the portfolio shares are equal at 0.46. In 2001, we have 0.62 in the PSID compared with 0.67 in the CEX. However, changes in log wealth show a peculiar pattern for the CEX, with Δw_{it} being negative for all years prior to 1995. This has to do with topcoding practices in the CEX. Before 1996, the critical value for topcoding of security holdings was set \$100,000. In 1996, this was raised to \$500,000, followed by further increases in later years. This further underscores the need for controls for and interactions with lagged wealth in our regressions, as well as for year dummies to control for the effects of changes in topcoding practices on the mean of Δw_{it}. As a robustness check, we also check our results based on the post-95 CEX sample. Finally, the bottom panel of Table 2 shows a slightly positive relationship between lagged wealth and the risky asset share, similar to the PSID.

[7] In principle, some of this measurement error could arise from topcoding, which leads to the loss of some observations of wealthy households who experienced further gains in wealth. Unlike in the CEX, however, the topcoding critical values in the PSID are very high (\$10 million for individual wealth components until the late 1990s and \$100 million subsequently).

4. Changes in Wealth and Asset Allocation

4.1 Regression specification

We start with our baseline regression, which is set up to measure the effect of changes in wealth on households' stock allocation. Recall from Section 2.2 that we estimate

$$\Delta\rho_{it} = \beta_0 + \beta_1 \Delta w_{it} + \beta_2 x_{it} + \varepsilon_{it} , \qquad (13)$$

where the subscript i denotes the household, $\Delta\rho_{it} = \rho_{it} - \rho_{it-1}$ is the change in the log of the percentage of financial wealth invested in stocks, Δw_{it} is the change in log total wealth, and x_{it} is a vector of household characteristics. We also explore a specification that employs consumption growth,

$$\Delta\rho_{it} = \beta_0 + \beta_1 \Delta c_{it} + \beta_2 x_{it} + \varepsilon_{it} , \qquad (14)$$

where Δc_{it} is the change in log consumption. In the PSID, the only available consumption information (and in years when the wealth questions where administered only from 1994 onwards) is food consumption, including outlays for food away from home. We calculate log consumption growth by comparing current food consumption to food consumption five years (two years in the last wave) earlier. Since food consumption is more likely to be smoothed than other more lumpy components of consumption, it may in fact be more closely related to the theoretical notion of consumption than broader definitions of consumption. On the other hand, of course, nonseparabilities with other consumption components and the fact that food is a necessity make it less suitable. In the CEX we have more detailed consumption information and we follow Parker (2001) and define nondurables consumption as the sum of food, alcohol, apparel, transportation, entertainment, personal care, and reading expenditure. We calculate semiannual consumption growth by comparing consumption reported at the time of the fourth and fifth interview with consumption at the second and third interview. We deseasonalize the CEX consumption growth data by regressing it on dummies for the interview month and using the residual in our tests.

Due to measurement problems, household measures of consumption growth are likely to have a low signal-to-noise ratio. We therefore also explore estimation of specification (14) with instrumental variables, as outlined in Section 2.2. We use growth in log wealth (Δw_{it}) and interactions of Δw_{it} with age and lagged wealth group dummies as instruments.

An important role of our controls in x_{it} is to eliminate life-cycle effects and effects that are common across individuals. We include dummies for age groups and the lagged quartile of wealth and we also interact them with Δw_{it}. In this way, the estimate of the coefficient β_1 picks up only the effects of wealth changes that are not driven by the stage of the life-cycle. Moreover, to avoid contamination of our results by spurious correlation from time trends in aggregate stock holdings, wealth, and consumption most of our specifications also employ year dummies.

4.2 Wealth regression results

Table 3 presents the results for the wealth regression, i.e., Equation (13), using data from the PSID. All specifications except the first include age and wealth group effects, as well as year dummies. Heteroskedasticity-consistent standard errors are reported in parentheses. Specifications (i) and (ii) differ only with respect to the inclusion of group and year effects. Controlling for life-cycle effects via the inclusion of age dummies apparently does not make much difference. The estimated coefficient on Δw_{it} is negative and about four standard errors below zero in both cases. Economically, though, the estimate is not far from zero: A coefficient of –0.17 implies that 10% growth in real wealth leads to a reduction of the share of risky assets by 1.7% (e.g., from 50% to 49.15%; recall that the coefficient represents an elasticity). Moreover, compared with the coefficient of 10-20—or, after allowing for attenuation due to measurement error, at least well above 1.0—that would be needed to generate sufficient time-variation in risk aversion to explain the "excess volatility" of stock returns, the estimate is clearly too small to indicate any economically significant wealth effects. This is also underscored by the low R^2.

To explore this further, in Specification (iii), we include the contemporaneous change in the value of real estate owned by the household ($\Delta Home_{it}$), gross of mortgages. The idea here is that a negative wealth effect might spuriously arise because households saving for the purchase of a home might do so with riskless assets, experiencing increasing wealth over time, but when the home is purchased eventually the holdings of riskless assets drop strongly (see, e.g., Faig and Shun 2002). The positive coefficient on $\Delta Home_{it}$ and the reduced magnitude of the coefficient on Δw_{it} is at least consistent with this concern. Nevertheless, the estimated wealth effect is still negative. In specification (iv) we interact Δw_{it} with dummies for age groups and lagged wealth quartiles. The results show that the estimated wealth effect on the risky asset share is about zero for all wealth and age groups except perhaps for the young and poor, for whom the estimated effect is positive, but still close to zero. Specification (v) includes a dummy *Trade* that takes a value of one if the household reported any net trade in risky assets from t-1 to t. We also interact this dummy with Δw_{it}. The estimated coefficient on this interaction term shows that non-trading households have a somewhat stronger negative wealth effect (-0.337) than those that did trade (-0.337+0.232 = -0.115).

One concern about these results—and one that we examine further in much more depth below—is that slow adjustment of households to wealth shocks might bias our coefficient estimate. For this reason, the dependent variable in our specification (vi) is $\Delta\rho_{it+1}+\Delta\rho_{it}$. This means that we allow an additional 2 years in the last wave, and 5 years otherwise, for the wealth effect on risky asset holdings to materialize. As the table shows, with –0.018 (std.err. 0.064) the magnitude of the estimated coefficient is smaller than before, suggesting some role for slow adjustment. Finally, in specifications (vii) and (viii) we also include equity in private businesses and farms in our risky asset holdings variable. As can be seen in the table, this does not make much difference. Overall, across all specifications, the estimated wealth effect is essentially about zero.

Table 4 repeats a set of similar regressions for the CEX sample. Recall that in the CEX, the first differences are calculated over annual intervals, i.e., over shorter periods than in the PSID. Nevertheless, we obtain similar results with the CEX, with estimates of roughly similar magnitude. For all regression specifications, the estimated wealth effect is negative. The only notable differences are that the coefficient on ΔHome is basically zero for the CEX[8], and that the young and poor have a more strongly negative wealth effect than wealthier households and households with older heads. The latter finding suggests that the slightly positive wealth effect found for the young and poor in the PSID may not be a robust phenomenon.

4.3 Consumption growth regression results

Table 5 presents estimates for regression specifications based on Equation (14), where we use consumption growth instead of Δw_{it}. Consumption growth is likely to be much more noisy than changes in wealth. On the other hand, transitory variation in wealth might not be reflected in consumption, which is a useful property. The first set of columns in the table shows results for the PSID, the second set shows estimates for the CEX. The number of observations for the PSID is lower than before because we have consumption growth data only for the three final waves of the survey. Age and wealth group effects and year dummies are included for all specifications and we report results both for OLS and 2SLS estimation.

Specification (i) shows that consumption growth is basically uncorrelated with Δc_{it}: The coefficient estimate of 0.003 is economically small and it is also much smaller than its standard error of 0.040. In specification (ii) we explore how the estimate differs across age and wealth groups. We find that young households are the only sub-group with a positive relationship between consumption growth and changes in the portfolio share of risky assets. For all other

[8] This may have to do with the fact that there seem to be few first-time homebuyers in the CEX, possibly because of difficulty in reaching such households (due to address changes) when sampling. In the (longitudinal) PSID, a lot of effort is put into tracking individuals and their offspring over time.

groups the total effect (which can be obtained by adding the coefficient on Δc_{it} and the interaction terms corresponding to the group) is essentially zero. The same conclusion applies in specification (iii) where the dependent variable is $\Delta\rho_{it+1}+\Delta\rho_{it}$ to allow for slow adjustment.

Of course, these coefficient estimates on Δc_{it} are likely to suffer from severe attenuation bias due to measurement noise. To the extent that measurement errors in Δc_{it} and Δw_{it} have little correlation, the two-stage least squares regression in specification (iv) should go some way in reducing the attenuation bias. In this regression, we obtain a coefficient of –1.391 (standard error 0.689). The fact that the estimate is negative should not come surprising, given that our instruments are Δw_{it} and its interactions with age and lagged wealth dummies. But it shows that if we use only the component in Δw_{it} that is related to consumption growth variations, we still obtain a negative wealth effect. As a note of caution, though, the instruments are quite weak in terms of statistical significance in the first stage, which means that second-stage estimates may be biased and standard errors understated (Staiger and Stock 1997).

The results for the CEX in specifications (v) to (vii) yield a similar picture. Using OLS, the estimated coefficient on Δc_{it} is close to zero. Two-stage least squares estimation, using similar instruments as with the PSID data, yields a negative coefficient of –0.225. For the CEX, where we have a larger sample than for the PSID, the relationships between the instruments and Δc_{it} in the first stage are estimated with somewhat higher precision than for the PSID. For several age and wealth groups, the estimated consumption-wealth elasticity in the first stage is positive and more than two standard errors above zero, while for others it is around zero.

4.4 Robustness checks

We have carried out a variety of robustness checks, which we briefly summarize here before exploring in more detail the effect of inertia. We replicated our regressions—with quantitatively similar results—without imposing the capital gains filter (see Sections 3.1 and 3.2), with the risky asset share defined over total wealth instead of financial wealth, and with net financial wealth (i.e., net of non-mortgage debt) in the denominator of the risky asset share. To check the effect of topcoding in the CEX, we ran our regressions for the post-1995 CEX subsample, which is less affected by topcoding problems. We have also run regressions including growth in real labor income for family head and wife as a proxy for growth in human capital, interacted with dummies for age and lagged wealth groups. Only for the lowest wealth group there is a slight positive relationship with $\Delta\rho_{it}$, but it is not statistically significant. For all other wealth groups, the estimated coefficient on labor income growth is negative, but it is not significantly different from zero in statistical terms. To check whether lumpy consumption expenditures, in particular for durables, might falsely lead to negative wealth shock observations, we explored regressions with CEX data, where we added total expenditure to time t wealth when computing Δw_{it}. The coefficient estimate on this modified Δw_{it} however is still negative. Finally, we explored whether the fact that our definition of the risky asset share cannot exceed one by construction might bias our results.[9] Yet, interacting Δw_{it} with dummies for the lagged share of risky assets, we find that the estimated wealth effect is negative for all levels of the past risky asset share. Since the issue should be more relevant for households with a risky asset share close to one, it therefore seems unlikely that the failure of our risky asset share variable to capture leverage has any effect on our results.

5. Inertia in Asset Allocation

As a potential explanation for the absence of the positive effect in our regressions, one might conjecture that inertia could play some role: Changes in wealth might typically accrue first in the form of liquid assets (e.g., on the checking account). If the household is slow to adjust its risky asset holdings to recent changes in wealth, for example, because of adjustment costs (transaction

[9] In the PSID, stock positions are reported net of margin loans, so we cannot assess the degree of leverage. Moreover, recall that we do not subtract debts from financial wealth in the denominator of the risky asset share.

costs or cognitive costs) or a belief that changes in wealth might be transitory, our previous regression using contemporaneous changes in wealth and risky asset holdings could be biased against finding a positive wealth effect.[10] Of course, if adjustment were slow, it would also affect the asset pricing implications—it would imply that conditional risk premia should react to wealth shocks with a lag—which is why we carefully examine and control for such inertia effects in this section. This will also shed some more light on what else determines individuals' asset allocations, if it is not changes in wealth.

5.1 Regression specification

We define a variable $\Delta\pi_{it}$ that is designed to capture the effects of inertia. It is the hypothetical change in $\Delta\rho_{it}$ that the household would have experienced between time t-1 and t under perfect inertia—that is, if it had not undertaken any purchases or sales of risky assets between t-1 and t. We then modify our wealth regression, Eq. (6), by including $\Delta\pi_{it}$:

$$\Delta\rho_{it} = \beta_0 + \beta_1 \Delta\pi_{it} + \beta_2 \Delta w_{it} + \beta_3 x_{it} + \varepsilon_{it} , \qquad (15)$$

If subjects exhibit perfect inertia with respect to their stock holdings (and $\beta_2 = 0$) then $\beta_1 = 1$. If households exhibit no inertia at all, and hence rebalance their portfolios immediately following capital gains and in- and outflows of financial wealth, then $\beta_1 = 0$. If households chase returns, in the sense that they buy more stocks following capital gains, then they exacerbate the effect of capital gains and it is possible that $\beta_1 > 1$, depending on their inertia with respect to in- and outflows.

Several identification issues are important to consider when interpreting regression (15). First, we can only estimate the average effect across our sample of households (except for interactions with demographic characteristics). It may well be the case that some households

[10] Of course, if inertia is mainly with respect to capital gains and capital gains drive a substantial portion of wealth changes, then inertia would induce a positive relationship between changes in wealth and changes in the risky asset share.

rebalance, some exhibit inertia, while others chase returns. Second, capital gains and losses, and hence $\Delta\pi_{it}$, are bound to be measured with large error. Classical measurement error in net purchases of risky assets (which is used to back out capital gains) would lead to attenuation bias for the β_1 estimate. More worrying would be underreporting of trades, i.e., a bias towards zero in reported trades. Especially in the PSID, where the recall period is up to five years, systematic forgetting of trades may be a frequent problem. This would lead to a spurious positive relationship between $\Delta\pi_{it}$ and $\Delta\rho_{it}$. Since the CEX is sampled at shorter frequencies than the PSID, the capital gains information there should be less prone to measurement error and thus, the CEX results will provide a useful check. Also, we can look at subsamples excluding the respondents reporting zero trades, which may be the most error-prone ones.

Mismeasurement of the level of stock holdings (which results in an error in $\Delta\pi_{it}$ of the same sign), may also create spurious positive correlation between $\Delta\rho_{it}$ and $\Delta\pi_{it}$ and an upward bias of the β_1 estimate. In the PSID, which is a true panel data set, we can address the latter type of measurement error with the following alternative specification:

$$\Delta\rho_{it+1} + \Delta\rho_{it} = \beta_0 + \beta_1\,\Delta\pi_{it} + \beta_2\,\Delta w_{it} + \beta_3\,x_{it} + \varepsilon_{it}\,, \tag{16}$$

Error in the measurement of the level of stock holdings should be uncorrelated across survey waves, implying that the difference in across two periods, i.e. $\Delta\rho_{it+1} + \Delta\rho_{it}$, should be unaffected by time t measurement error. If inertia is perfect (and returns on the household's portfolio are unpredictable), then we should still find $\beta_1 = 1$.

Apart from these identification issues, regression (16) is also interesting in economic terms. If inertia results from adjustment costs that are traded off against the benefits of rebalancing, then there should inertia in the short run, but in the long run, the share of stock holdings should revert to its optimal level. Since the measurement periods in the PSID are five years (except for the last wave, where it is two years), allowing an additional period for rebalancing to take place should eliminate most of the inertia effects if they are driven by adjustment costs.

Finally, while regression (15) is suitable for estimation of wealth effects controlling for inertia, it cannot distinguish between inertia with respect to capital gains and inertia with respect to other in- and outflows to and from financial wealth. For that reason we also investigate a specification that separates out the effect of capital gains related inertia on the stock portfolio share. Here, the dependent variable is the growth in log dollar stock holdings, denoted Δs_{it}:

$$\Delta s_{it} = \beta_0 + \beta_1 r_{it} + \beta_2 \Delta f_{it} + \beta_3 \Delta w_{it} + \beta_3 x_{it} + \varepsilon_{it} \qquad (17)$$

In this specification, r_{it} is the capital gain (i.e., the household's realized risky asset return excluding dividends) from t-1 to t and Δf_t is the growth in log financial wealth (gross of non-mortgage debt) excluding the effect of capital gains. Disregarding measurement issues for a moment, perfect inertia with respect to capital gains in this setting would imply $\beta_1 = 0$, returns-chasing $\beta_1 > 0$, and complete rebalancing of capital gains shocks would imply $\beta_1 = -1$. Perfect inertia with respect to in- and outflows of financial wealth (which are assumed to accrue initially as cash) would imply $\beta_2 = 0$ and no inertia at all would imply $\beta_2 = 1$. In this setting, underreporting of trades would lead to a spurious positive correlation between r_{it} and Δs_{it}, and so would mis-measurement in the level of stock holdings.

5.2 Results of inertia tests

Table 6 presents the results of our inertia tests, both for the PSID and CEX. Specifications (i), (ii), and (iii) show that inertia is strong. Our estimate for the coefficient on $\Delta\pi_{it}$ in specification (i) is 0.842, which is remarkably close to one. Compared with the earlier wealth regressions, the R^2, now around to 80%, has increased dramatically. In specification (ii) we find that the young and poor show somewhat less inertia than other groups, but even for them our estimate is still 0.558. To explore to what extent underreporting of trades might drive the results, we interact $\Delta\pi_{it}$ with our Trade dummy in specification (iii). Not surprisingly, the coefficient estimate for those reporting zero trades is almost exactly equal to one, but, more interestingly, for those who did trade, the implied inertia coefficient is still close to 0.8. Since part of the concern about measurement error is caused by the relative frequency of zero trades in the PSID, these results are reassuring.

Most importantly, specifications (i) to (iii) show that the estimated wealth effect is still very close to zero. The coefficient on Δw_{it} varies from 0.031 to 0.051, with marginal statistical significance. Compared with our earlier results in Table 3 where the effect was slightly negative, this suggests that slow adjustment to in- and outflows does play some role in generating the slightly negative relationship between Δw_{it} and $\Delta \rho_{it}$ that we found in Table 3.

In specification (iv) we can see that the effect of inertia is very persistent. The coefficient estimate of 0.599 shows that even after five years (two years for the final survey wave), much of the inertia effect still persists. This casts some doubt on a pure adjustment cost story for this inertia. After all, many households did undertake trades between t-1 and t+1, they just did not undo the effects of capital gains or losses. In specification (iv) we also find a somewhat larger positive wealth effect (0.159, standard error 0.057). Yet, even allowing for several additional years to adjust, and controlling for initial inertia, the effect is still at least an order of magnitude too small compared with the values delivered by our calibration exercise in Section 2. When the reaction to wealth shocks is delayed by several years is also likely to be difficult to generate counter-cyclical variation in risk premia from wealth shocks.

From specification (vi), we can see that using the CEX data produces a somewhat smaller estimate of the coefficient on $\Delta \pi_{it}$ (0.555, standard error 0.046). The reason might be that the shorter recall period in the CEX reduces the measurement error in capital gains and eliminates some spurious positive correlation between $\Delta \pi_{it}$ and $\Delta \rho_{it}$. For the CEX, too, the estimated wealth effect is basically zero (-0.006, standard error 0.014). In summary, these results show that after controlling for inertia, changes in wealth have virtually no impact on changes in asset allocation—which is what one would expect, for example, if agents had constant relative risk aversion preferences (CRRA preferences would not explain inertia, though). These results strengthen our conclusion from the previous section that time-varying risk aversion at the aggregate level does not seem to be explained by wealth effects at the household level.

In specifications (v) and (vii) we further explore the origins of inertia. Recall that in these regressions, perfect inertia would imply coefficients of zero on capital gains (r_{it}) and changes in financial wealth (Δf_{it}). Table 6 shows that with respect to capital gains, there is indeed almost perfect inertia. The coefficient estimates of –0.017 (standard error 0.019) and –0.156 (standard error 0.027) for the PSID and CEX, respectively, imply that households hardly do any rebalancing following capital gains and losses. In contrast, the coefficient estimates on Δf_{it}, which are 0.320 (PSID) and 0.438 (CEX), show that households at least partly adjust their portfolio in response to in- and outflows of financial wealth, although they still go less than half the way that perfect rebalancing would require. Thus, the most important source of inertia seems to be the failure of households to rebalance following capital gains and losses. This implies that capital gains and losses on stocks have a special influence on a household's asset allocation. At a given point in time, it depends to a large extent on its own history of capital gains and losses. Moreover, given the extant evidence that households' are poorly diversified (see, e.g., Goetzmann and Kumar 2004), these capital gains and losses are largely idiosyncratic and thus, inertia generates cross-sectional variation in risk asset shares.

6. Discussion

The absence of a positive effect of wealth shocks on risky asset holdings in our micro-data is inconsistent with the idea that individual investors become less risk averse in booms *because* they experienced increases in wealth, and that they are more risk averse in recessions, *because* their wealth has declined. Since such a wealth effect does not seem to exist at the household level, it cannot serve as a micro-foundation for representative-agent models with counter-cyclical variation in risk aversion. Of course, this does not rule out that these models could be rescued by differently specified micro-foundations. To do so, several avenues seem possible.

First, our analysis of household portfolios has focused only on a subset of the investor universe. Perhaps it is institutional investors rather than individuals whose risk aversion varies over time. Alas, for institutions it is not clear that the standard utility-over-consumption approach is appropriate in the first place. With delegated portfolio management, agency and contracting issues may be important, and could also influence how risk aversion varies over time. Second, from Constantinides and Duffie (1996) we know that with a judiciously specified labor income process and incomplete markets, we can generate any desired marginal utility process for a representative agent.

This underscores that it is important to study micro-implications of asset-pricing models, as the same representative-agent model may lead to very different predictions, depending on its micro-foundation.[11] For example, compare a world with habit-formation individuals with a Constantinides-Duffie world of idiosyncratic income risk. What moves risk premia in the first model are wealth shocks. What moves risk premia in the second world is time-variation in the cross-sectional dispersion of income shocks. While it is possible, in theory, to construct the same marginal utility process for a representative agent from both approaches, the two stories could lead to drastically different predictions about and how when risk premia should change.

Our finding that capital gains and losses have a large and highly persistent impact on asset allocation is difficult to reconcile with preferences that do not give special roles to different components of wealth. For example, in the standard habit formation model, changes in wealth should affect marginal utility in the same way, irrespective of their origin from capital gains, say, or from labor income. One explanation could be that individuals are not willing to rebalance their portfolios because they perceive it as too costly, or because they are biased towards the status-quo.

[11] Lettau (2001) evaluates the Constantinides-Duffie model with income data from the PSID and argues that idiosyncratic income risk does not help much to explain asset pricing puzzles.

Similar inertia has been documented for investor behavior in 401(k) retirement accounts. Samuelson and Zeckhauser (1988), Ameriks and Zeldes (2001), Agnew, Balduzzi, and Sunden (2000), and Huberman and Sengmueller (2004) find that a large portion of individuals hardly ever trade at all in their retirement accounts, and that inflow allocations are rarely changed.[12]

At least a part of the phenomenon could also arise from preferences or beliefs. What appears, at the surface, to be inertia might in fact be a deliberate asset allocation choice by households. For example, in Barberis, Huang, and Santos (2001), narrow framing of risks and dynamic first-order risk aversion imply that it is recent capital gains, rather than changes in total wealth, that make investors less averse to bearing stock market risk. Alternatively, capital gains might make investors more optimistic about future stock returns. Our data does not allow us to disentangle the effect of past and contemporaneous stock returns, but some suggestive evidence has been obtained in other settings. Benartzi (2001) and Huberman and Sengmueller (2004) find that individuals allocate more of their retirement assets towards company stock if it has recently performed well. Vissing-Jorgensen (2003) provides survey evidence on investor expectations suggesting that individual investors during the late 90's boom expected future stock returns to be high, not low, which could be the result of extrapolative expectations. Our findings on the special role of past capital gains and losses in asset allocation suggest that an exploration of these avenues may be promising to find an explanation of the equity volatility puzzle that is consistent with investor behavior at the micro-level.

[12] Our results differ from Odean (1998), though, but they are not inconsistent with his findings. Odean shows that investors tend to hold on to *individual stocks* that are losers and that they tend to sell winners. Our results, in contrast, concern the allocation to the *entire portfolio of stocks*, including mutual funds investments, and they also incorporate the effect of purchases, not just sales.

References

Agnew, Julie, Pierluigi Balduzzi, and Annika Sunden, 2004, Portfolio Choice and Returns in a Large 401(k) Plan, *American Economic Review*, forthcoming.

Ameriks, John, and Stephen P. Zeldes, 2001, How Do Household Portfolio Shares Vary With Age?, working paper, Columbia University.

Bakshi, Gurdip S., and Zhiwu Chen, 1996, The Spirit of Capitalism and Stock-Market Prices, *American Economic Review* 86, 133-157.

Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect Theory and Asset Prices, *Quarterly Journal of Economics* 116, 1-53.

Barksy, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, 1997, Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study, *Quarterly Journal of Economics* 112, 537-579.

Benartzi, Shlomo, 2001, Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock, *Journal of Finance* 56, 1747-64.

Brav, Alon, George M. Constantinides, and Christopher C. Geczy, 2002, Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence, *Journal of Political Economy* 110, 793-824.

Campbell, John Y., 2000, Asset Pricing at the Millenium, *Journal of Finance* 55, 1515-1567.

Campbell, John Y., 1993, Intertemporal Asset Pricing Without Consumption Data, *American Economic Review* 83, 487-512.

Campbell, John Y., and John H. Cochrane, 1999, By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior, *Journal of Political Economy* 107, 205-251.

Campbell, John Y., and Albert S. Kyle, 1993, Smart Money, Noise Trading, and Stock Price Behavior, *Review of Economic Studies* 60, 1-34.

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Inflation Illusion and Stock Prices, *American Economic Review Papers and Proceedings* 94, 19-23.

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 2000, Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?, *American Economic Review* 90, 787-805.

Constantinides, George M., and Darrell Duffie, 1996, Asset Pricing with Heterogeneous Consumers, *Journal of Political Economy*, 104, 219-240.

Dynan, Karen E., 2000, Habit Formation in Consumer Preferences: Evidence from Panel Data, *American Economic Review*, 90, 391-406.

Faig, Miquel, and Pauline Shun, 2002, Portfolio Choice in the Presence of Personal Illiquid Projects, *Journal of Finance* 57, 303-328.

Goetzmann, William N., and Alok Kumar, 2004, Diversification Decisions of Individual Investors and Asset Prices, working paper, Yale University.

Guiso, Luigi, Michael Haliassos, and Tullio Jappelli, 2003, Household Stockholding in Europe: Where Do We Stand and Where Do We Go?, *Economic Policy* 18, 523-577.

Juster, F. Thomas, Joseph P. Lupton, James P. Smith, and Frank Stafford, 2004, The Decline in Household Saving and the Wealth Effect, working paper, University of Michigan.

Hansen, Lars Peter, and Kenneth J. Singleton, 1993, Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns, *Journal of Political Economy* 91, 249-265.

Heaton, John, and Deborah Lucas, 2000, Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk, *Journal of Finance* 55, 1163-1198.

Huberman, Gur, and Paul Sengmueller, 2004, Performance Predicts Asset Allocation: Company Stock in 401(k) Plans, *Review of Finance* 8, 403-443.

Lettau, Martin, 2001, Idiosyncratic Risk and Volatility Bounds, or, Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?, *Review of Economics and Statistics*, forthcoming.

Lupton, Joseph P., 2003, Household Portfolio Choice and the Habit Liability: Evidence from Panel Data, working paper, Federal Reserve Board.

Mankiw, N. Gregory, 1986, The Equity Premium and the Concentration of Aggregate Shocks, *Journal of Financial Economics*, 17, 211-219.

Mankiw, N. Gregory, and Stephen Zeldes, 1991, The Consumption of Stockholders and Nonstockholders, *Journal of Financial Economics* 29, 97-112.

Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *Journal of Finance* 53, 1887-1934.

Parker, Jonathan A., 2001, The Consumption Risk of the Stock Market, *Brookings Papers on Economic Activity*, 279-348.

Routledge, Bryan R., and Stanley E. Zin, 2003, Generalized Disappointment Aversion and Asset Prices, working paper, NBER.

Samuelson, William, and Richard Zeckhauser, 1998, Status Quo Bias in Decision Making, *Journal of Risk and Uncertainty* 1, 7-59.

Staiger, Douglas, and James H. Stock, 1997, Instrumental Variables Regression with Weak Instruments, *Econometrica* 65, 557-586.

Vissing-Jorgensen, Annette, 2002, Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure, working paper, Northwestern University.

Vissing-Jorgensen, Annette, 2003, Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions, *NBER Macroeconomics Annual* 2003.

Vissing-Jorgensen, Annette, and Orazio P. Attanasio, 2003, Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion, *American Economic Review Papers and Proceedings* 93, 384-391.

TABLE 1. SUMMARY STATISTICS, PSID

	All Family Units			Stock market participants in t-1					
	#Obs.	%risky	W_{it}	#Obs.	%risky	W_{it}	r_{it}	Δw_{it}	$\Delta \rho_{it}$
				By Year					
1989	2443	0.12	166	311	0.46	407	0.39	0.39	0.08
		(0.25)	(513)		(0.31)	(752)	(0.98)	(0.72)	(1.11)
1994	2667	0.18	190	459	0.59	427	0.57	0.27	0.42
		(0.32)	(578)		(0.30)	(772)	(1.04)	(0.69)	(1.13)
1999	2340	0.23	263	544	0.64	643	0.52	0.43	0.26
		(0.34)	(1148)		(0.27)	(2060)	(1.06)	(0.85)	(0.97)
2001	3113	0.20	233	742	0.62	590	0.09	0.18	0.03
		(0.33)	(1140)		(0.28)	(2124)	(0.87)	(0.59)	(0.80)
				By Lagged Wealth Quartile					
1st		0.05	21		0.55	99	0.49	0.66	0.19
		(0.18)	(132)		(0.29)	(102)	(1.07)	(0.93)	(1.02)
2nd		0.12	57		0.57	230	0.38	0.35	0.18
		(0.26)	(111)		(0.30)	(204)	(1.02)	(0.58)	(1.05)
3rd		0.20	143		0.58	398	0.36	0.21	0.16
		(0.31)	(338)		(0.29)	(248)	(0.93)	(0.52)	(0.95)
4th		0.36	631		0.61	1,339	0.34	0.06	0.27
		(0.36)	(1543)		(0.28)	(2653)	(0.92)	(0.60)	(0.95)

Note: We calculate %risky as the percentage of financial wealth invested in stocks, and $\Delta \rho_{it}$ denotes the first difference in its log. W_{it} denotes total household wealth, r_{it} is the log capital gain on household i's stock portfolio over the last five years (two years in 2001), and Δw_{it} is the first difference of log wealth. Wealth data are deflated by the CPI into 2001 dollars. The first panel shows cross-sectional means with standard deviations in parentheses, the second panel shows wave-by-wave cross-sectional means and standard deviations for lagged wealth quartiles averaged across years. The first set of columns considers the full sample, the second set of columns shows summary statistics for the sub-sample of households with nonzero stock holdings in the prior survey wave.

TABLE 2. SUMMARY STATISTICS, CEX

	All family units			Stock market participants in t-1					
	#Obs.	%risky	W_{it}	#Obs.	%risky	W_{it}	r_{it}	Δw_{it}	$\Delta\rho_{it}$
				By Year					
1988	1255	0.10	84	234	0.46	138	-0.03	-0.18	-0.05
		(0.24)	(98)		(0.31)	(110)	(0.25)	(0.77)	(0.40)
1989	1346	0.11	82	252	0.46	148	0.05	-0.08	-0.01
		(0.25)	(102)		(0.32)	(124)	(0.20)	(0.40)	(0.37)
1990	1351	0.10	82	206	0.51	146	0.04	-0.14	-0.01
		(0.24)	(104)		(0.30)	(127)	(0.21)	(0.44)	(0.32)
1991	1190	0.12	93	233	0.49	146	0.02	-0.02	0.01
		(0.26)	(99)		(0.31)	(115)	(0.23)	(0.44)	(0.26)
1992	1402	0.11	86	253	0.48	141	0.02	-0.06	0.03
		(0.25)	(100)		(0.31)	(134)	(0.23)	(0.63)	(0.33)
1993	1310	0.10	82	217	0.50	140	0.02	-0.07	0.02
		(0.24)	(94)		(0.30)	(122)	(0.24)	(0.68)	(0.39)
1994	1422	0.13	80	260	0.53	129	0.00	-0.02	0.01
		(0.27)	(92)		(0.31)	(109)	(0.20)	(0.50)	(0.32)
1995	1284	0.12	73	233	0.52	127	0.04	0.00	0.05
		(0.26)	(84)		(0.31)	(108)	(0.20)	(0.44)	(0.32)
1996	543	0.18	92	115	0.69	189	0.08	0.09	0.01
		(0.33)	(127)		(0.29)	(201)	(0.27)	(0.48)	(0.20)
1997	1169	0.18	92	253	0.64	183	0.09	0.00	0.02
		(0.33)	(120)		(0.30)	(160)	(0.24)	(0.41)	(0.21)
1998	1192	0.20	96	276	0.71	191	0.07	0.04	0.01
		(0.35)	(131)		(0.28)	(184)	(0.23)	(0.49)	(0.36)
1999	1161	0.20	104	263	0.69	213	0.05	0.05	0.03
		(0.35)	(131)		(0.29)	(212)	(0.21)	(0.32)	(0.16)
2000	1427	0.20	103	339	0.69	204	0.06	-0.01	0.02
		(0.35)	(146)		(0.28)	(195)	(0.26)	(0.39)	(0.33)
2001	1363	0.19	91	296	0.67	172	-0.11	-0.02	-0.01
		(0.34)	(1245)		(0.29)	(165)	(0.31)	(0.36)	(0.24)
				By Lagged Wealth Quartile					
1st		0.04	7		0.54	32	0.06	0.10	0.00
		(0.15)	(23)		(0.31)	(35)	(0.27)	(0.73)	(0.38)
2nd		0.10	37		0.57	91	0.03	-0.06	0.02
		(0.25)	(35)		(0.31)	(48)	(0.22)	(0.40)	(0.32)
3rd		0.15	85		0.57	173	0.02	-0.06	0.01
		(0.29)	(47)		(0.30)	(57)	(0.23)	(0.34)	(0.24)
4th		0.29	226		0.61	351	0.00	-0.09	0.02
		(0.35)	(131)		(0.27)	(146)	(0.21)	(0.28)	(0.20)

Note: We calculate %risky as the percentage of financial wealth invested in stocks, and $\Delta\rho_{it}$ denotes the first difference in its log. W_{it} denotes total household wealth, r_{it} is the log capital gain on household i's stock portfolio over the last five years (two years in 2001), and Δw_{it} is the first difference of log wealth. Wealth data are deflated by the CPI into 2001 dollars. The first panel shows cross-sectional means with standard deviations in parentheses, the second panel shows wave-by-wave cross-sectional means and standard deviations for lagged wealth quartiles averaged across years. The first set of columns considers the full sample, the second set of columns shows summary statistics for the sub-sample of households with nonzero stock holdings in the prior survey wave.

TABLE 3. WEALTH EFFECTS IN ASSET ALLOCATION, PSID

Model	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Dependent Variable	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it+1}+\Delta\rho_{it}$	$\Delta\rho^{Bus}_{it}$	$\Delta\rho^{Bus}_{it}$
Δw_{it}	-0.145	-0.173	-0.071	0.152	-0.337	-0.018	-0.058	0.002
	(0.040)	(0.042)	(0.050)	(0.073)	(0.074)	(0.064)	(0.037)	(0.044)
ΔChildren	0.029	0.012	0.008	0.019	0.005	0.055	0.025	0.022
	(0.040)	(0.040)	(0.040)	(0.040)	(0.040)	(0.057)	(0.035)	(0.035)
ΔHome			0.165					0.097
			(0.051)					(0.046)
Trade					0.118			
					(0.050)			
$Age_{(30\ to\ 50)} \times \Delta w_{it}$				-0.296				
				(0.097)				
$Age_{(50\ to\ 70)} \times \Delta w_{it}$				-0.008				
				(0.123)				
$Age_{(>70)} \times \Delta w_{it}$				-0.202				
				(0.148)				
$w_{it-1(Quart.2)} \times \Delta w_{it}$				-0.206				
				(0.120)				
$w_{it-1(Quart.3)} \times \Delta w_{it}$				-0.062				
				(0.127)				
$w_{it-1(Quart.4)} \times \Delta w_{it}$				-0.303				
				(0.114)				
$Trade \times \Delta w_{it}$					0.232			
					(0.086)			
Intercept	0.228	0.386	0.305	0.129	0.287	0.467	0.311	0.263
	(0.025)	(0.128)	(0.127)	(0.136)	(0.130)	(0.288)	(0.125)	(0.124)
Age group fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wealth group fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	1%	3%	4%	4%	5%	2%	2%	2%
#Obs.	2,023	2,023	2,023	2,023	2,023	907	2,023	2,023

Note: The dependent variable $\Delta\rho_{it}$ is the first difference in the log of a household's share of risky assets in its financial wealth portfolio. In columns (vii) and (viii), equity in private business is included in risky asset holdings. The independent variable of main interest is the change in log real wealth (Δw_{it}). The control variables include the change in the number of children in the household (ΔChildren), the change in the total value (gross of mortgages) of real estate owned by the household (ΔHome), a dummy that equals one if the household sold or bought risky assets between t and t-1 (Trade), and interactions with lagged wealth quartile, age group, and trade dummies. Only households with nonzero stock holdings in the lagged survey wave are included in the regressions. Estimation is by pooled OLS. Heteroskedasticity-consistent standard errors are reported in parentheses.

TABLE 4. WEALTH EFFECTS IN ASSET ALLOCATION, CEX

	(i)	(ii)	(iii)	(iv)	(v)
Δw_{it}	-0.053	-0.056	-0.065	-0.223	-0.077
	(0.022)	(0.022)	(0.037)	(0.103)	(0.031)
ΔHome			-0.004		
			(0.021)		
Trade					0.057
					(0.012)
$Age_{(30\ to\ 50)} \times \Delta w_{it}$				0.175	
				(0.098)	
$Age_{(50\ to\ 70)} \times \Delta w_{it}$				0.141	
				(0.108)	
$Age_{(>70)} \times \Delta w_{it}$				0.235	
				(0.099)	
$w_{it-1(Quart.2)} \times \Delta w_{it}$				0.048	
				(0.053)	
$w_{it-1(Quart.3)} \times \Delta w_{it}$				0.063	
				(0.042)	
$w_{it-1(Quart.4)} \times \Delta w_{it}$				0.042	
				(0.046)	
Trade $\times \Delta w_{it}$					0.055
					(0.040)
Intercept	0.009	-0.043	-0.037	-0.036	-0.062
	(0.045)	(0.043)	(0.046)	(0.040)	(0.043)
Age group effects	No	Yes	Yes	Yes	Yes
Wealth group effects	No	Yes	Yes	Yes	Yes
Year effects	No	Yes	Yes	Yes	Yes
Adj. R^2	1%	1%	1%	2%	2%
#Obs.	3,339	3,339	3,222	3,339	3,339

Note: The dependent variable $\Delta\rho_{it}$ is the first difference in the log of a household's share of risky assets in its financial wealth portfolio. The independent variable of main interest is the change in log real wealth (Δw_{it}). The control variables include the change in the total value (gross of mortgages) of real estate owned by the household (ΔHome), a dummy that equals one if the household sold or bought risky assets between t and t-1 (Trade), and interactions with lagged wealth quartile, age group, and trade dummies. Only households with nonzero stock holdings in the lagged survey wave are included in the regressions. Estimation is by pooled OLS. Heteroskedasticity-consistent standard errors are reported in parentheses.

TABLE 5. WEALTH EFFECTS IN ASSET ALLOCATION: REGRESSIONS WITH HOUSEHOLD CONSUMPTION GROWTH, PSID AND CEX

	PSID				CEX		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
Method	OLS	OLS	OLS	2SLS	OLS	OLS	2SLS
Dependent Variable	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it+1} + \Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$
Δc_{it}	0.003	0.336	0.026	-1.391	-0.008	0.113	-0.225
	(0.040)	(0.088)	(0.084)	(0.689)	(0.022)	(0.072)	(0.240)
ΔChildren	-0.012	-0.012	0.085	0.11			
	(0.050)	(0.050)	(0.067)	(0.091)			
$Age_{(30\ to\ 50)} \times \Delta c_{it}$		-0.369				-0.163	
		(0.134)				(0.068)	
$Age_{(50\ to\ 70)} \times \Delta c_{it}$		-0.302				-0.037	
		(0.153)				(0.086)	
$Age_{(>70)} \times \Delta c_{it}$		-0.349				-0.073	
		(0.155)				(0.063)	
$w_{it-1(Quart.2)} \times \Delta c_{it}$		0.011				-0.057	
		(0.158)				(0.084)	
$w_{it-1(Quart.3)} \times \Delta c_{it}$		-0.032				0.014	
		(0.132)				(0.084)	
$w_{it-1(Quart.4)} \times \Delta c_{it}$		-0.02				-0.072	
		(0.138)				(0.088)	
Intercept	0.444	0.386	0.401	0.696	-0.042	-0.045	-0.042
	(0.116)	(0.110)	(0.637)	(0.246)	(0.042)	(0.042)	(0.045)
Age group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wealth group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	N/A	Yes	Yes	Yes	Yes
Adj. R^2	2%	2%	2%	0%	1%	1%	0%
#Obs.	1250	1250	408	1244	3319	3319	3234

Note: The dependent variable $\Delta\rho_{it}$ is the first difference in the log of a household's share of risky assets in its financial wealth portfolio. The independent variable of main interest is the change in log real consumption (Δc_{it}). The control variables include the change in the number of children in the household (ΔChildren), and interactions with lagged wealth quartile and age group dummies. Only households with nonzero stock holdings in the lagged survey wave are included in the regressions. Estimation is by pooled OLS, except for columns (iv) and (vii), which show results of two-stage least squares estimation. In the latter case, Δc_{it} is instrumented with Δw_{it} (change in log real wealth) and interactions of Δw_{it} with age-group and wealth-quartile dummies. Heteroskedasticity-consistent standard errors are reported in parentheses.

TABLE 6. INERTIA IN ASSET ALLOCATION, PSID AND CEX

	PSID					CEX	
Model	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
Dependent variable	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it}$	$\Delta\rho_{it+1} + \Delta\rho_{it}$	Δs_{it}	$\Delta\rho_{it}$	Δs_{it}
$\Delta\pi_{it}$	0.842	0.558	1.005	0.599		0.555	
	(0.022)	(0.056)	(0.005)	(0.042)		(0.046)	
Δw_{it}	0.042	0.051	0.031	0.159		-0.006	
	(0.023)	(0.023)	(0.022)	(0.057)		(0.014)	
r_{it}					-0.017		-0.156
					(0.019)		(0.027)
Δf_{it}					0.320		0.438
					(0.026)		(0.049)
ΔChildren	-0.008	-0.009	-0.016	0.007	0.015		
	(0.019)	(0.019)	(0.018)	(0.047)	(0.031)		
Trade			0.248				
			(0.017)				
$Age_{(30\ to\ 50)} \times \Delta\pi_{it}$		0.258					
		(0.064)					
$Age_{(50\ to\ 70)} \times \Delta\pi_{it}$		0.221					
		(0.074)					
$Age_{(>70)} \times \Delta\pi_{it}$		0.167					
		(0.110)					
$w_{it-1(Quart.2)} \times \Delta\pi_{it}$		0.036					
		(0.062)					
$w_{it-1(Quart.3)} \times \Delta\pi_{it}$		0.067					
		(0.062)					
$w_{it-1(Quart.4)} \times \Delta\pi_{it}$		0.118					
		(0.058)					
Trade $\times \Delta\pi_{it}$			-0.227				
			(0.030)				
Intercept	0.465	0.411	0.287	0.399	0.320	0.048	0.018
	(0.084)	(0.074)	(0.071)	(0.324)	(0.308)	(0.025)	(0.025)
Age group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wealth group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	78%	78%	80%	32%	21%	34%	22%
#Obs.	1,900	1,900	1,900	850	1,910	3,332	3,423

Note: The dependent variable is $\Delta\rho_{it}$, the first difference in the log of a household's share of risky assets in its financial wealth portfolio, and, in columns (iv) and (vi), Δs_{it}, the first difference in the log of dollar risky asset holdings. The independent variables of main interest are $\Delta\pi_{it}$ the hypothetical value of $\Delta\rho_{it}$ that would have prevailed without any rebalancing on part of the household and with all in- and outflows of financial wealth accruing to risky asset holdings, Δw_{it}, the change in log real wealth (Δw_{it}), r_{it}, the log capital gain on the household's risky asset holdings, and Δf_{it}, the first difference in log financial wealth. The control variables include the change in the number of children in the household (ΔChildren), a dummy that equals one if the household sold or bought risky assets between t and t-1 (Trade), and interactions with lagged wealth quartile, age group, and trade dummies. Only households with nonzero stock holdings in the lagged survey wave are included in the regressions. Estimation is by pooled OLS. Heteroskedasticity-consistent standard errors are reported in parentheses.