

NBER WORKING PAPER SERIES

DO WEALTH FLUCTUATIONS GENERATE TIME-VARYING RISK AVERSION?
MICRO-EVIDENCE ON INDIVIDUALS' ASSET ALLOCATION

Markus K. Brunnermeier
Stefan Nagel

Working Paper 12809
<http://www.nber.org/papers/w12809>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2006

We thank John Campbell, Darrell Duffie, Mark Gertler, Francisco Gomes, Joy Ishii, Frank de Jong, Christian Julliard, Martin Lettau, Chris Malloy, Filippas Papanikolaou, Jonathan Parker, Jacob Sagi, Ken Singleton, Ilya Strebulaev, Annette Vissing-Jorgensen, Yihong Xia, Motohiro Yogo, two anonymous referees, and seminar participants at the CEPR Meetings in Gerzensee, the Five-Star Conference at NYU, HECER Helsinki, Humboldt University Berlin, IAEEG Trier, London School of Economics, the Stanford-Berkeley joint Finance seminar, and UC Irvine for useful comments. Brunnermeier acknowledges financial support from the National Science Foundation and the Alfred P. Sloan Foundation. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2006 by Markus K. Brunnermeier and Stefan Nagel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Markus K. Brunnermeier and Stefan Nagel

NBER Working Paper No. 12809

December 2006

JEL No. G11

ABSTRACT

We use data from the PSID to investigate how households' portfolio allocations change in response to wealth fluctuations. Persistent habits, consumption commitments, and subsistence levels can generate time-varying risk aversion with the consequence that when the level of liquid wealth changes, the proportion a household invests in risky assets should also change in the same direction. In contrast, our analysis shows that the share of liquid assets that households invest in risky assets is not affected by wealth changes. Instead, one of the major drivers of households' portfolio allocation seems to be inertia: households rebalance only very slowly following inflows and outflows or capital gains and losses.

Markus K. Brunnermeier
Princeton University
Department of Economics
Bendheim Center for Finance
Princeton, NJ 08540
and NBER
markus@princeton.edu

Stefan Nagel
Stanford University
Graduate School of Business
518 Memorial Way
Stanford, CA 94305
and NBER
nagel_stefan@gsb.stanford.edu

1 Introduction

A growing number of studies in macroeconomics and finance propose models in which agents' relative risk aversion is time-varying. The most popular approach is to use habit-formation preferences, in particular difference habits, which imply that felicity is a function of consumption minus a habit. In asset pricing, difference-habit models have some success in reproducing the mean and counter-cyclical of asset return risk premia found in the data (Constantinides (1990); Bakshi and Chen (1996); Campbell and Cochrane (1999)). In macroeconomics, habits help to jointly match stylized facts about asset returns and the business cycle (see, e.g., Jermann (1998); Boldrin, Christiano, and Fisher (2001)). An alternative approach focuses on consumption commitments, which can have effects similar to those of difference habits, in particular, similar time-variation in relative risk aversion (Chetty and Szeidl (2005)).

While habit preferences¹ seem to help in matching aggregate data, little is known yet about whether the predictions of habit-formation models also fit with microdata. Mehra and Prescott (2003), for example, point out that it is not clear whether investors actually have the huge time varying counter-cyclical variations in risk aversion postulated by models like Campbell and Cochrane (1999). One of the key implications of difference habits is that individuals' relative risk aversion should vary with wealth, in contrast to models with constant relative risk aversion (CRRA). An increase in wealth, for example, should lead to a temporary decrease in relative risk aversion. This is an important, but so far untested prediction. In this paper, we provide evidence on this question from microdata on how households allocate their wealth between risky and riskless assets.

To clarify the implications of difference habits for asset allocation, we start by studying a simple discrete-time model of portfolio choice. The issues are most transparent if we take the view that with CRRA preferences—i.e., without habit—the investor would have sufficiently low risk aversion so that she would invest most of her liquid wealth in risky assets. If we now introduce a difference habit, this increases the desire to hold riskless assets. Their primary role is to provide sufficient financial resources to ensure that future consumption can always be kept above the level of the habit. Hence, optimal riskless asset holdings are tied to the slow-moving habit level and thus relatively fixed. But liquid wealth fluctuates, due to capital gains, income, and consumption. As a result, when liquid wealth increases, the optimal share of risky assets in the liquid wealth portfolio increases, and vice versa. Effectively, relative risk aversion varies with wealth.

We test this prediction with household-level panel data from the Panel Study of

¹In the following we often speak, for ease of reference, somewhat loosely of "habit preferences" or "habit formation", but we mean difference habits (or subsistence levels, or consumption commitments that lead to similar effects), which lead to time-varying risk aversion. But we exclude ratio habits of the type used by Abel (1990), because they imply constant relative risk aversion.

Income Dynamics (PSID), covering a period of about 20 years. We first examine how changes in liquid wealth affect stock market participation. We find that changes in liquid wealth have a significant positive effect on the probability of stock market entry and a negative effect on the probability of exit. While this is consistent with time-varying risk aversion if there are some fixed per-period cost of participation, similar effects also arise with CRRA preferences. Thus, these tests cannot discriminate between habit models and models with CRRA preferences.

Unlike for the participation decision, we find that changes in liquid wealth essentially play no role in explaining changes in asset allocation for households that participate in the stock market. We regress the change in proportion of liquid assets invested in risky assets on the change in liquid wealth and find that the positive effect predicted by difference-habit models is absent. If anything, the effect is slightly negative (but economically tiny). This is not the result of low statistical power—our coefficients are quite precisely estimated. Thus, the asset allocation results favor the CRRA model.

Our regressions control for a broad set of household characteristics, including variables related to the life-cycle and time dummies to eliminate aggregate effects and focus on cross-sectional variation. We also pay attention to measurement error. We obtain similar results when we instrument changes in wealth with independently measured income growth and inheritances, albeit with somewhat lower precision. Moreover, we also show, theoretically, that it doesn't matter whether the liquid wealth change is anticipated, as long as the anticipated change is not entirely riskless. What matters is that optimal riskless asset holdings are relatively fixed in the short-run, because they are tied to the habit level, and thus any fluctuation in current liquid wealth, whether previously anticipated or not, leads to a change in the risky asset share.

One possible explanation for the lack of a contemporaneous effect of wealth changes on asset allocation is that households' asset allocation is governed by inertia. When capital gains and losses arise, they are not rebalanced, and when in- and outflows arise, they affect mostly the riskless asset (cash) balances. With infrequent or delayed adjustment, the first effect would lead to a positive, the latter effect to a negative relationship between changes in liquid wealth and the liquid risky asset share. Indeed, we find that inertia seems to be the dominant factor determining changes in asset allocation. The PSID data on purchases and sales of risky assets allows us to reconstruct, approximately, how the portfolio allocation would look like if households had not bought or sold risky assets between successive interview dates (assuming that all in- and outflows affect only cash balances). We find that actual portfolio allocations are quite close. The data on purchases and sales are surely noisy and probably affected by forgotten trades, but the strength of the inertia effect seems to be too big to be just the result of measurement error.

Given that there seems to be strong inertia, we then check whether a positive effect of liquid wealth changes on portfolio shares might appear if we allow for slow adjustment. We regress future changes in the risky asset share on past changes in

wealth and find a small positive effect. But in terms of economic magnitudes it is again a very small effect, and it is statistically weak.

Taken together, our findings suggest that relative risk aversion does not vary with wealth changes in the way postulated by habit-formation models. The large variations in relative risk aversion induced by wealth changes that these theories predict are evidently absent from microdata. At least with respect to the relationship between asset allocation and wealth, our evidence suggests that constant relative risk aversion is a good description of microeconomic behavior. But the CRRA model cannot explain the large inertia in households' portfolio shares either.

Our evidence on household asset allocation ties in well with some recent work that finds it hard to reconcile habit preferences and microdata along other dimensions of households' economic choices. Dynan (2000) finds no evidence that household-level consumption growth exhibits the patterns predicted by internal habit-formation models. Gomes and Michaelides (2003) study a life-cycle model of consumption and portfolio choice and find that the introduction of habit formation makes it more difficult to match empirical regularities in microdata. A recent paper by Sahm (2006) examines relative risk aversion measures elicited from responses to hypothetical gamble questions in the Health and Retirement Study and finds no effect of wealth changes on changes in relative risk aversion. The findings in these studies contrast with Lupton (2003) who finds a negative relationship between past consumption levels and current risky asset holdings, including businesses and real estate, which he argues is consistent with habit formation, and Ravina (2005), who finds support for habit formation in credit card purchases data. The results in our (first-differences) regressions are also consistent with earlier evidence that the cross-sectional relationship between the level of the risky asset share (Heaton and Lucas (2000); Guiso, Haliassos, and Jappelli (2003)) or elicited relative risk aversion measures (Barksy, Juster, Kimball, and Shapiro (1997)) and the level of wealth is essentially flat among households that participate in the stock market.

The paper is organized as follows. Section 2 presents a simple portfolio choice model with habit preferences, our methodology, and the data. Section 3 reports our main results. In Section 4 we discuss the implications of our results.

2 Theory and Methodology

2.1 Model of Asset Allocation with Habits

We develop a simple model of portfolio choice that illustrates how relative risk aversion can be time-varying when agents' preferences exhibit difference habits, subsistence levels, or similar features. Let time be discrete and consider a single agent with infinite horizon. The agent's wealth at time t is denoted W_t and is measured before time t consumption, C_t . There are two securities the agent can invest in: a risky asset, with return R_t and a riskfree asset with constant return R_f . At time t the agent chooses C_t

and the proportion of $W_t - C_t$ invested in the risky asset, α_t , to solve the problem

$$\max E_t \sum_{\tau=0}^{\infty} \delta^\tau \frac{(C_{t+\tau} - X)^{1-\gamma}}{1-\gamma}, \quad (1)$$

subject to the intertemporal budget constraint

$$W_{t+1} = (1 + R_{p,t+1})(W_t - C_t), \quad (2)$$

where δ is the subjective discount factor, γ is the curvature of the felicity function, $R_{p,t+1} \equiv \alpha_t(R_t - R_f) + R_f$ is the return on the investors' liquid wealth portfolio, and X is the habit. Consumption paths with $C_t \leq X$ at some future t with non-zero probability are assigned infinitely negative utility. We assume that risky asset returns have a log-normal distribution. Because we focus on cross-sectional differences between households and not on aggregate variation, we also assume, for simplicity, constant expected returns and constant volatility.

In our basic discussion we assume that X is constant. This should be thought of as an approximation to a model in which X varies slowly. In the appendix we show that our basic model can be viewed as an approximation to an internal habit model along the lines of Constantinides (1990), where habit responds sluggishly to past consumption. As Campbell and Cochrane (1999) point out, letting the habit level respond slowly, over several years, to changes in consumption is necessary to match empirical features of asset returns, such as a highly persistent price-dividend ratio, persistent volatility, and long-run forecastability of returns, and so we focus on the effects of such slow moving habits in our analysis. Alternatively, X could represent an external habit that does not depend on the action of our single agent, a constant subsistence level, or the cash-flow stream required to finance future committed consumption along the lines of Chetty and Szeidl (2005).

We solve the agent's problem by redefining consumption and wealth such that the objective and the budget constraint map into a standard CRRA problem, for which we know the relevant properties of the solution. Define surplus wealth $W_t^* \equiv W_t - \frac{X}{R_f} - X$ and surplus consumption $C_t^* \equiv C_t - X$. We can rewrite the maximization problem as

$$\max E_t \sum_{\tau=0}^{\infty} \delta^\tau \frac{C_{t+\tau}^{*1-\gamma}}{1-\gamma}. \quad (3)$$

Now assume that the investor at time t invests a fraction α_t^* of $W_t - C_t - \frac{X}{R_f}$ into the risk asset, and the rest in the riskless asset. This surplus portfolio yields a return $R_{p,t+1}^* \equiv \alpha_t^*(R_t - R_f) + R_f$. The remaining $\frac{X}{R_f}$ dollars are invested in the riskless asset. Without further restrictions on α_t^* , this decomposition of the wealth portfolio into two components is without loss of generality. The budget constraint now becomes

$$W_{t+1} = (1 + R_{p,t+1}^*) \left(W_t - C_t - \frac{X}{R_f} \right) + (1 + R_f) \frac{X}{R_f}, \quad (4)$$

from which we obtain

$$W_{t+1}^* = (1 + R_{p,t+1}^*) (W_t^* - C_t^*). \quad (5)$$

Thus, our problem maps into the problem of a CRRA investor with wealth W_t^* , consumption C_t^* , and risky asset portfolio share α_t^* . If expected returns and volatility are constant, $\alpha_t^* = \alpha^*$, i.e., it is constant, as we know from Samuelson (1969). Then the risky asset share of our habit utility investor, as a fraction of post-consumption wealth at time t , is

$$\alpha_t = \alpha^* \left(1 - \frac{X}{(W_t - C_t) R_f} \right). \quad (6)$$

We now see that α_t is increasing in W_t , holding X constant. The agent invests the present value of the future habit, X/R_f , in riskless assets, and surplus wealth over and above that amount like a CRRA investor. Hence, if W_t is close to X/R_f , the agent's effective relative risk aversion is high, because most of the wealth is used to self-insure the stream of future habit. If W_t is a lot larger than X/R_f , the agent invests approximately like a CRRA investor.

So far, the agent's wealth is composed entirely liquid wealth, by which we mean the sum of stocks and riskless bonds. In a more realistic model, however, the household would also have so-called background wealth, i.e., labor income, housing wealth, and perhaps wealth in a private business. Labor income and business ownership are sources of risky income. Home equity represents a risky asset if the household expects to sell at some point, say at retirement (or trade to a house of different size, or in a different location).² The presence of background wealth complicates the relationship between liquid wealth levels, habit, and the proportion of liquid wealth allocated to stocks, because α^* , and hence also α_t , depend on how much background wealth the agent has, and on the riskiness of background wealth. This makes the portfolio choice problem quite untractable.

Fortunately, the existing literature shows that as long as the returns on background wealth have a relatively low correlation with stock returns (for which there is some empirical support), the main effect of risky background wealth is a diversification effect. That is, the presence of background wealth allows the household to diversify away some of the risks of stocks, reducing the aversion to holding stocks.³ This diversification effect can simplify the problem.

²Costs of adjusting the holdings of illiquid background wealth can also lead to an additional effect of house ownership that relates to one of the possible interpretations of X . Recognizing that housing is often financed by collateralized borrowing (mortgage), and that many households own a house for most of their life, one can view housing as a long-lived durable good to which the household has committed financial resources of future periods, and whose consumption is very costly to adjust. In this case, the mortgage payments have the same consequences as a habit. The agent uses riskless assets to self-insure and be able to meet future mortgage payments to avoid large adjustment costs of trading down to a smaller house. This effect is analyzed in Chetty and Szeidl (2005) and X can be viewed as incorporating the effects of habits and of such committed future payments.

³It's important to distinguish here between models that examine the effect of adding mean-zero

The very purpose of habit-formation models is to try to explain facts about asset returns and consumption with moderate values for the felicity curvature parameter γ . For example, Campbell and Cochrane (1999) set $\gamma = 2$ and the habit is responsible for the lion's share of effective relative risk aversion and its variation over time; the curvature of the felicity function would induce very little relative risk aversion in the absence of the habit. For such moderate values of γ , realistically calibrated models of household portfolio choice with CRRA preferences and background wealth, such as Bertaut and Haliassos (1995), Heaton and Lucas (1997), Cocco (2004), Cocco, Maenhout, and Gomes (2005), and Yao and Zhang (2005) all find that households should basically invest 100% of their liquid wealth in stocks. In other words, $\alpha^* = 1$. Note also that 100% is the upper bound (i.e., no leverage) for the allocation to risky assets in discrete time if background wealth is risky (in the sense that there is a strictly positive probability that background wealth can fall to an arbitrarily small value over the next period).

These results suggest that we can incorporate the effects of housing wealth and labor income into our model by approximating $\alpha^* \approx 1$, i.e., by assuming that a CRRA investor without habit would invest about 100% of his liquid wealth in stocks. By making this approximation, we effectively assume that the habit does most of the lifting required to get risky asset shares below 100% in our portfolio choice model. This is eminently consistent with the spirit of habit-formation equilibrium models where the habit does most of the lifting to get a sizeable equity risk premium.

The benefit of this approximation is that the optimal portfolio share no longer depends on background wealth and that it varies over time only because of variation in $\frac{X}{W_t - C_t}$, not because of variation in α^* :

$$\alpha_t \approx 1 - \frac{X}{(W_t - C_t) R_f}. \quad (7)$$

Eq. (7) provides the basis for the main tests in the paper. The analyses where we assume $\alpha^* \approx 1$ should be interpreted as a test of the joint hypothesis that habits lead to time-variation in risky asset shares and that the variation induced by habits dominates the effects of variation in background risks.

But we also examine alternative specifications where we relax the $\alpha^* \approx 1$ assumption. Allowing for a range of values that are still relatively close to one, say $0.9 \leq \alpha^* \leq 1$, would make little difference as long as variation in $\frac{X}{W_t - C_t}$ dominates variation in α^* . If α^* is instead substantially smaller than one, we have to take into account that a changing composition of wealth can change the degree to which stock market risk is diversified away in background wealth. We relax the $\alpha^* \approx 1$ approximation in two heuristic ways.

background risks (e.g., Gollier and Pratt (1996)) from the portfolio choice literature that's our focus here, where background *wealth* is added, which has positive mean returns.

First, when we look at the liquid risky asset share, we control for the relative magnitude of background wealth. We use the labor income/liquid wealth ratio interacted with age, the business wealth/liquid wealth ratio, and the housing wealth/liquid wealth ratio, which should help capture variation in α^* . In the end, we find that these controls have little effect. This is consistent with cross-sectional analyses in the literature. Heaton and Lucas (2000) and Yao and Zhang (2005) find little evidence that these asset composition ratios are correlated with liquid risky asset shares (among households that participate in the stock market), perhaps with the exception of the relative share of business wealth, which seems to have some small negative impact. These empirical findings also support our approximation assumption $\alpha^* \approx 1$.

Second, we look at the financial wealth (liquid plus housing plus business wealth) risky asset share. On the right-hand side of Eq. (6) one then also has to redefine W_t as financial wealth. The only background wealth component left is human wealth, which is difficult to measure. We include the labor income/financial wealth ratio interacted with age as a proxy for background human wealth.

2.2 Implications for Time-variation in Risky Asset Holdings

We are interested in the implications of habit formation for time-variation in the willingness to hold risky assets. To obtain our estimating equation, we first linearize Equation (7),

$$\begin{aligned}\alpha_t &= 1 - \frac{X}{(W_t - C_t) R_f} \\ &= 1 - \exp(x - w_t) \\ &\approx \kappa - \rho(x - w_t).\end{aligned}\tag{8}$$

where $x \equiv \log(X/R_f)$, $w_t \equiv \log(W_t - C_t)$ ⁴, and the last approximate equality follows from a first-order Taylor approximation, where ρ and κ are constants, and $\rho > 0$. Intuitively, when W_t is close to X , changes in W_t have a big effect on X/W_t at the margin, but the effect is small when W_t is big. This feature is preserved by linearizing around the log habit-wealth ratio, such that α_t is linear in log liquid wealth, so that the bigger W_t , the smaller the marginal impact of an increase in W_t .

Taking first differences of equation (8), we get

$$\Delta\alpha_t = \rho\Delta w_t.\tag{9}$$

This result depends on our assumption that X is approximately constant. As discussed above, it can be justified, for example, by looking at idiosyncratic wealth shocks and X

⁴Note that in our empirical data, we measure post-consumption wealth each period, so the definition of w_t corresponds to the definition of wealth in the data.

as an external habit that is not affected by idiosyncratic wealth shocks, or with X as an internal habit that is slowly moving, and so reacts only very sluggishly to changes in wealth. Eq. (9) forms the basis for our empirical tests.

Anticipated vs. unexpected wealth changes—When taking our model to the data, an important question is whether the relationship between changes in liquid wealth and portfolio shares would be modified if changes in liquid wealth are partly anticipated. It turns out that this distinction does not matter for the validity of our tests.

Consider the following example. At $t = 1$, the agent expects to receive a big one-time payment, for example an inheritance from a rich uncle, in period $t = 2$. Let's assume that the probability of getting the inheritance is high, and the risk of obtaining or not obtaining it (for example, because the uncle prefers to donate the money to charity) and the value of the uncle's assets are uncorrelated with stock market returns. The key point is that a small probability of not getting the inheritance is sufficient to make the anticipated inheritance unsuitable as a means to insure future habit. Hence, at $t = 1$, the agent still needs to invest $\frac{X}{R_f}$ in riskless assets to insure future habit, despite the anticipated inheritance. Only when the inheritance is actually received at $t = 2$, but not before, liquid wealth increases relative to $\frac{X}{R_f}$, and the risky asset share increases.

In this example, the dollar amount of stock holdings moves one-for-one with realized changes in the dollar amount of liquid wealth, despite the fact that the agent anticipated the change in liquid wealth. One can show that a similar logic applies when the agent invests less than 100% of surplus wealth in stocks, and with non-constant, but slow-moving habits, and when one incorporates the consumption-savings decision. In Appendix A.2 we numerically solve a three-period model that illustrates these effects.⁵

In summary, it is not crucial for our tests to distinguish between anticipated and unanticipated changes in liquid wealth. What we need to control for, however, is that portfolio shares of risky assets and wealth may have some common predictable life-cycle pattern, for reasons unrelated to habit. But this is a different issue that we address below.

Stock market participation—In the model above, the agent would always participate in the stock market ($\alpha_t > 0$), because the optimal investment policy ensures that $W_t - C_t > \frac{X}{R_f}$ (given sufficient initial wealth W_0). However, if one extended the model

⁵It may be useful at this point to draw an analogy to the consumption literature. Power utility implies a precautionary savings motive. In that case, consumers faced with an expectation of a big inflow in the near future, but which can be low or zero with strictly positive probability would hesitate to run down savings to raise today's consumption, because that would expose them to a small, but significant risk that consumption might be extremely low in the future if the anticipated inflow does not realize (Zeldes (1989), Carroll (1997)). In other words, consumption displays "excess" sensitivity to anticipated changes in income. In the same way, an investor with habit preferences faced with an anticipated inflow is deterred from increasing the risky asset share by a possibly very small probability that this inflow might not realize. As a consequence, the portfolio share does not react until the income is realized, as long as there is some residual uncertainty about this income change.

and allowed for some costs of participating in the stock market, the household might choose non-participation (see, e.g., Vissing-Jorgensen (2002), Gomes and Michaelides (2005)). These costs might be of financial nature, they could be opportunity costs of time and attention, and also psychological costs.

Suppose one extended the model to include a fixed per-period participation costs. Then, changes in liquid wealth could induce stock market entry or exit. For example, a household experiencing a negative change in wealth might choose to exit. This happens for two reasons. The first is that with a lower amount of wealth, the benefits from investing in stocks are smaller relative to the fixed level of costs. This effect would arise even with CRRA preferences. Second, in the model with habit, as liquid wealth declines, the agent wants to invest a smaller amount of the shrunken liquid wealth in stocks, further reducing the benefits from participating.

The latter effect suggests that the presence of habits could lead to time-varying stock market participation. We therefore also look at the empirical relationship between changes in liquid wealth and stock market entry and exit. But it's important to keep in mind that this is only a first step. A finding that stock market participation varies with changes in liquid wealth will, on its own, not be sufficient to discriminate between CRRA and the habit model.

2.3 Econometric Issues

In our data, we have observations on wealth and asset holdings at dates that are $k = 5$ years or $k = 2$ years apart, depending on the subsample. Therefore, we rewrite Eq. (9) as $\Delta_k \alpha_t = \rho \Delta_k w_t$, where Δ_k denotes a k -period first-difference operator, $\Delta_k y_t \equiv y_t - y_{t-k}$.

To arrive at our estimating equation, we must take into account that variables outside of the model may cause common movements in the level of liquid assets and the risky asset share. For example, it is possible that both α_t and w_t have some correlated deterministic pattern over the life-cycle. Therefore, we condition on a vector q_{t-k} of household characteristics that should capture such patterns, if present. It includes variables that are either constant or known at $t - k$, and a vector of ones. In addition, most of our Specifications also include $\Delta_k h_t$, a vector of variables that capture major changes in family composition or asset ownership that could lead to preference shifts that are possibly correlated with $\Delta_k w_t$. Finally, we add a mean-zero error term ε_t , which captures unobserved forces on the portfolio share that are outside of the model and that are uncorrelated with $\Delta_k w_t$, q_{t-k} , and $\Delta_k h_t$. Thus, our estimating equation is

$$\Delta_k \alpha_t = \beta q_{t-k} + \gamma \Delta_k h_t + \rho \Delta_k w_t + \varepsilon_t. \quad (10)$$

To reduce clutter, we continue to omit household subscripts. In our basic specification, we assume that w_t is well-measured so that ε_t is uncorrelated with $\Delta_k w_t$ and the other

regressors and we can estimate Eq. (10) with ordinary least squares (OLS). Below we also consider violations of this assumption when w_t and α_t are measured with error.

Life-cycle effects and preference shifters—Our conditioning variables in q_{t-k} include a broad range of variables related to the life-cycle, background, and the financial situation of the household at $t - k$. For lack of a better name, we refer to them collectively as life-cycle controls. We include age and age²; indicators for completed high school and college education, respectively, and their interaction with age and age²; dummy variables for gender and their interaction with age and age², marital status, health status; the number of children in the household, the number of people in the household; dummy variables for any unemployment in the k years leading up to and including year $t - k$, and for coverage of the household head’s job by a union contract. In addition, we include the log of the equity in vehicles owned by the household, log family income at $t - k - 4$, 2-year growth in log family income at $t - k$ and $t - k - 2$, and a variable for inheritances received in the k years leading up to and including year $t - k$. The other category of control variables, labelled as preference shifters, $\Delta_k h_t$, includes changes in some household characteristics between $t - k$ and t : changes in family size, changes in the number of children, and a sets of dummies for house ownership, business ownership, and non-zero labor income at t and $t - k$. The idea behind the house ownership dummies, for example, is that households might save for the purchase of a home with mostly 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 dummies at t and $t - k$ should absorb those effects.

Idiosyncratic vs. aggregate wealth changes—Our partial equilibrium portfolio choice model deals with the decision of a single household, holding constant aggregate quantities and prices. But if a wealth change is common to all households, and hence they all want to change their exposure to risky assets, the effect on asset allocation is dampened, because instead of quantity, it’s now the price (and thus, the expected return) of risky assets that adjusts. To uncover the effects of habits, we must therefore eliminate aggregate changes in wealth and asset holdings and focus on household-specific variation. For this reason, we include time fixed effects in q_{it-k} , which effectively de-means wealth changes and risky asset shares cross-sectionally. In Appendix A.3 we show that our estimator is consistent as the number of cross-sectional units $N \rightarrow \infty$. In addition, we also recognize that there could be local effects, where asset holdings and household income and other sources of wealth variation are tied to the local economy. To eliminate such local effects as much as possible, we interact the year dummies with dummies for the four PSID geographical regions, which provides us with a set of year-region dummies.

Measurement error—Measurement error is a standard concern with microdata from surveys. We model measured wealth as $\tilde{w}_t \equiv w_t + u_t$, i.e. the sum of true wealth and a measurement error u_t , which implies $\Delta_k \tilde{w}_t = \Delta_k w_t + \Delta_k u_t$. We assume that measurement error is uncorrelated with true wealth. More precisely, we assume that

$Cov(u_{t+i}, w_t) = 0$ for $i = -1, 0, 1$, so that $Cov(\Delta_k u_t, \Delta_k w_t) = 0$. Let the measured risky asset share be $\tilde{\alpha}_t = \alpha_t + v_t$, with measurement error v_t . Substituting wealth and the risky asset share into Eq. (10), we obtain

$$\Delta_k \tilde{\alpha}_t = \beta q_{t-k} + \gamma \Delta_k h_t + \rho \Delta_k \tilde{w}_t + \varepsilon_t + v_t - \rho \Delta_k u_t \quad (11)$$

Measurement error renders OLS inconsistent because $\Delta_k \tilde{w}_t$ and the composite residual $\varepsilon_t + v_t - \beta \Delta_k u_t$ are correlated. First, $Cov(\Delta_k \tilde{w}_t, \Delta_k u_t) > 0$, which biases the coefficient estimate for ρ towards zero. Second, because the numerator (stocks) and denominator (stocks plus riskless assets) of α_t are made up by components of w_t , we should also expect that $Cov(\Delta_k \tilde{w}_t, v_t) \neq 0$, which could also bias the sign of the coefficient. Specifically, measurement error only in stock holdings would lead to $Cov(\Delta_k \tilde{w}_t, v_t) > 0$, measurement error only in riskless assets would lead to $Cov(\Delta_k \tilde{w}_t, v_t) < 0$. It is not possible to unambiguously sign the combined effect when both stock holdings and riskless assets are mismeasured. Overall, the bias in the estimate of ρ could go in either direction, depending on whether measurement error in stocks or riskless assets dominates.

To address the measurement error problem, we look for instrumental variables for $\Delta_k w_t$. The identification requirement is that the instruments, z_t , are (partially) correlated with $\Delta_k w_t$, but not with $\varepsilon_t + v_t - \beta \Delta_k u_t$. Given such instruments, we can estimate ρ consistently with two-stage least squares (TSLS). Our instruments are quantile dummies for income growth from $t - k$ to t (similar to Dynan (2000) in a different application), and inheritance receipts (as in Meer, Miller, and Rosen (2003)) between $t - k$ and t . These instruments are based upon survey questions that are different from those for the components of w_t . Hence, it is reasonable to assume that the elements of z_t are uncorrelated with $\varepsilon_t + v_t - \beta \Delta_k u_t$.

Unfortunately, it is to be expected that we lose precision compared with the OLS estimator, and so there is a tradeoff between potential measurement-error bias and precision. A priori, it is not clear that the TSLS estimator will be closer to the true parameter in a mean-squared error sense. Therefore, we report both the OLS and TSLS results in our tests.

2.4 Data

We use data from the Panel Study of Income Dynamics (PSID), obtained from the University of Michigan. It is a longitudinal study that tracks family units and their offspring over time. We use data on asset holdings collected in the years 1984, 1989, 1994, 1999, and 2001, and 2003. Income data and many household characteristics are available annually until 1997, and every second year from then onwards. Appendix A.4 describes our data in more detail. Here we briefly discuss the definition of the variables that we extract from the files. To make magnitudes comparable over time, we deflate

all income and wealth data by the consumer price index (CPI) into December 2001 dollars.

Variable definitions—We define liquid assets as the sum of holdings of stocks and mutual funds plus riskless assets, where we follow common practice and define riskless assets as the sum of cash-like assets and holdings of bonds. Subtracting other debts, which comprises non-mortgage debt such as credit card debt and consumer loans, from liquid assets yields liquid wealth. We further denote the sum of liquid wealth, equity in a private business, and home equity as financial wealth.

We then calculate two risky asset shares: First, the sum of stocks and mutual funds held, divided by liquid assets (the liquid risky asset share).⁶ Second, the sum of stocks and mutual funds, home equity, and equity in a private business, divided by financial wealth (the financial risky asset share).

As our income variable we use total family income. The inheritance variable included in the vector of household characteristics q_{t-k} is the value of inheritances received scaled by income to adjust for the fact that a given amount of inheritance has different relevance for households with different income and wealth. More precisely, we measure it as the log of one plus the value of the inheritance divided by family income at $t - k - 4$, in the $k = 5$ subsample, and $t - k - 2$ in the $k = 2$ subsample.⁷ We want to use this variable as an instrument for liquid wealth, so we obviously cannot scale by wealth, and we scale by income instead.

In years when the wealth questions were administered, the PSID asked subjects to report on the amount of stocks and mutual funds 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). This information allows us to decompose the change in the amount of stocks and mutual funds held into an active investment/disinvestment component and a capital gains/losses component (see Appendix A.4 for details). There is reason to expect that this active investment information is noisy and that some households may systematically forget trades. We do not use the capital gains and active investment information in our main tests, only later when we examine inertia effects. We discuss the measurement error issue at that point.

Sample selection and weighting—To be included in our sample, we require that the marital status of the family unit head remained unchanged from $t - k$ to t and that no assets were moved in or out as a consequence of a family member moving into or out of the family unit. We also exclude observations on households if the household

⁶Since we don't have more detailed information on the composition of risky asset holdings, we don't know, for example, whether households hold stocks with low or high systematic risk or how well they are diversified. However, our focus is on whether households reallocate between stocks and riskless assets in response to wealth changes, not on whether they reallocate between stocks of different systematic risk (which the theory is silent about).

⁷We don't use $t - k - 5$ income in the earlier subsample, because we want to avoid using income data from prior to survey year 1980 when topcoding of income was much more prevalent than in later years.

head is retired at t . For risky asset shares to be meaningful, we also require a certain minimum level of wealth. We therefore exclude households with liquid wealth less than \$10,000 or financial wealth less than \$10,000 at $t - k$. Overall, the data requirements are quite demanding. In many of our regressions we need observations on income at $t - k - 2$, and $t - k - 4$. This means that we need households that participate in many consecutive waves of the survey. We weight observations with PSID sample weights when we present summary statistics, but we do not use the sample weights in our regression analyses, because doing so would be inefficient (Deaton (1997), p. 70). In any case, as Appendix A.5 shows, weighted regressions produce similar results.

Summary statistics—Table 1 presents summary statistics. The two top panels show pooled cross-section/time-series statistics for all households that satisfied the data and minimum lagged wealth requirements to be included in the sample. The two bottom panels show statistics for stock market participants, that is, those households that have stock holdings greater than zero at t and $t - k$. We further report separate summary statistics for our 1984-1999 sample, for which the time-span between successive waves of the PSID with wealth information is $k = 5$ years, and the 1999-2003 sample, for which $k = 2$.

As the table shows, the proportion of households participating in the stock market is 45% in the 1984-1999 sample, and 58% in 1999-2003. The large fraction of non-participants and the upward trend over time is roughly consistent with previous studies (e.g., Vissing-Jorgensen (2002)), but here the participation rate is somewhat higher because we focus on households that satisfy our minimum wealth requirements. The stock market entry variable in the two top panels is a dummy that is set to one for households that did not participate at $t - k$ and participate at t , and zero if the household does not hold stocks in $t - k$ and t . For households that participated in $t - k$, the variable is set to missing (therefore the lower number of observations). The stock market exit variable is defined in similar manner. It is equal to one for participants in $t - k$, but not t , zero for those that participated in $t - k$ and t , and missing otherwise. The numbers in the table show that there is considerable turnover in the group of participants. On average, between 34-35% of non-participants at $t - k$ enter the stock market until t , while about 19-24% of participants choose to exit. In our first tests below, we explore whether the probability of entry and exit is related to changes in liquid wealth.

Comparing wealth and income means and medians for all households and those for stock market participants, it is apparent that stock market participants have higher wealth and income on average. Combined with the fact that much of aggregate wealth is concentrated at the top end of the wealth distribution, wealthy households are, in some respects, the most important group of stock holders. Because extremely wealthy households have low response rates in surveys, they are not well represented in the PSID.⁸ However, Juster, Smith, and Stafford (1999) find that the wealth data in the

⁸An additional concern is topcoding, where observations above a certain threshold are set to the

PSID lines up well with data from the Survey of Consumer Finances (SCF) (which oversamples high income households and provides better data at the top end of the wealth distribution, but does not have a panel structure) at least up through the 98th percentile of the wealth distribution. Hence, our data should give us a good picture of the asset allocation choices of wealthy households, except for the extremely wealthy. But it is also useful to keep in mind that for testing the habit formation theory it is not crucial to have data from the very top end of the wealth distribution, because the theory does not predict that households with moderate levels of wealth should behave differently from households with very high levels of wealth.

The distribution of wealth and income has strong positive skewness. But when we examine changes in wealth and income, we difference logs. As the two bottom panels show, taking logs eliminates much of the skewness. The distribution of k -period differences in log wealth and log income is roughly symmetric. The 10th and 90th percentiles show that the k -period changes in log wealth are substantial, in particular in the 1984-1999 period, where $k = 5$.

The habit formation model predicts that these changes in wealth should give rise to changes in the risky asset share. As the statistics for the proportions of liquid and financial wealth invested in risky assets (%liq. assets risky and %fin. wealth risky) show, there is large variation in these risky asset shares over time. Whether these changes are related to wealth fluctuations is the subject of our main tests.

3 Results

3.1 Wealth Changes and Stock Market Participation

We start by investigating how changes in liquid wealth relate to stock market participation. There is existing evidence that higher wealth is associated with a higher probability that a household participates in the stock market (Bertaut and Haliassos (1995); Mankiw and Zeldes (1991); Vissing-Jorgensen (2002)), but this evidence is cross-sectional and does not necessarily speak to the dynamic relationship between changes in wealth and entry and exit. It is also possible that levels of liquid wealth are correlated with some unobserved fixed household characteristics that cause participation. Differencing removes the effect of these household characteristics.

Table 2 presents the results of probit regressions. In the first two columns, we estimate the probability of a household that did not participate at $t - k$ to enter the stock market until time t . In columns three and four we estimate the probability that a household that is participating at $t - k$ exits the stock market until t . The table shows

threshold value to protect the identity of the household. But in the PSID these cutoffs are very high (\$10 million per wealth component until the late 1990s and \$100 million subsequently) and affect only a very small number of cases.

the marginal effects, that is, the effect on the probability of entry or exit, evaluated at the sample means of the explanatory variables. The regressions include all the preference shifters and life-cycle controls we mentioned in Section 2.3. The focus of our interest is on the coefficient for the change in log liquid wealth.

As the table shows, in both samples (1984-1999 and 1999-2003) we find a positive coefficient, with high statistical significance. The point estimate of 0.124 in the first column implies that an increase in liquid wealth by 10% implies a roughly 1% increase in the probability to participate in the stock market. Hence, it is not a large effect, but it's not negligible either. The exit regressions in columns three and four show that the probability of exiting the stock market is negatively related to changes in liquid wealth. The magnitudes of the point estimates are a little smaller than for the entry regressions, but they, too, are different from zero at a high level of statistical significance. That changes in liquid wealth are significantly related to stock market entry and exit also provides some reassurance on the measurement error issue. Evidently, measured changes in liquid wealth are not driven entirely by measurement noise, otherwise we wouldn't find a significant relationship with stock market participation.

In summary, changes in liquid wealth appear to be one of the factors that causes changes in stock market participation. The reliably positive effect we find is consistent with time-varying risk aversion due to wealth changes, but it is also consistent with CRRA preferences in a model with fixed per-period participation costs.

3.2 Wealth Changes and Asset Allocation

We now turn to our main tests, looking at changes in the risky asset share conditional on participation, i.e., for those households that participate in the stock market at $t - k$ and t . Our goal is to estimate Eq. (10), and we do so with OLS and TSLS.

First Stage—Table 3 presents the TSLS first-stage estimates. The instruments are two indicator variables for log income growth between $t - k$ and t below the 10th or above the 90th percentile (see Table 1 for the value of these percentiles). Furthermore, we include an instrument for whether the household reports to have received an inheritance between $t - k$ and t (see Section 2.4 for definition).

The results in the table show that the instruments have a significant partial correlation with changes in log liquid wealth (columns 1 and 2) and changes in log financial wealth (columns 3 and 4) and the directions of the estimated effects are reasonable: higher income growth and a higher inheritance are associated with higher growth in liquid and financial wealth. The partial R^2 of the instruments is between 0.01 and 0.02, which suggests that the instruments still leave a large fraction of variation in wealth changes unexplained. This is typical for microdata. Nevertheless, the instruments are jointly highly significant, with p -values smaller than 0.005 for each of the Specifications. The F -statistics are, however, a bit lower than the rule of thumb of 10 suggested by Staiger and Stock (1997), below which the TSLS estimator is likely to

have some bias towards the OLS estimator and size-distorted confidence intervals due to weak instruments.⁹ In Appendix A.5 we re-estimate our regressions with methods that are robust for weak instruments along the lines of Moreira (2003) and find similar point estimates, albeit with wider confidence intervals.

Changes in liquid risky asset shares—Table 4 presents our main results. We regress changes in the liquid risky asset share on changes in liquid wealth. The habit model predicts that we should find a positive coefficient, but as the table shows the point estimates are very close to zero. In fact, for the OLS estimate in column 1 for the 1984–1999 sample we can reject at conventional significance levels that the coefficient is greater than zero. However, economically the estimate is basically zero. The coefficient of -0.013 in column 1 implies that 10% growth in real wealth leads to a tiny reduction in the share of risky liquid assets by 0.0013, e.g. from 50% to 49.87%. For the 1999–2003 sample, the estimate in column 4 is slightly positive, but again of tiny magnitude and statistically not significantly different from zero. The low explanatory power of wealth changes is also underscored by the low R^2 in these regressions, where essentially none of the variables, including the controls, explains an economically significant portion of changes in risky asset shares.

Having the two subsamples is useful, because they differ in the length of time between wealth measurement points. If habits are not sufficiently sluggish in catching up with consumption, having $k = 5$ years could be too long in the sense that there would be relatively quick mean reversion in risky asset shares, and so our regressions might not pick up much of the correlation with wealth changes. However, in the 1999–2003 subsample, we have $k = 2$ years, and we still find coefficient estimates that are virtually zero economically. It is, of course, still possible that we miss habit effects on the risky asset shares at even higher frequencies—but such high-frequency effects cannot be those that drive the slow-moving variation in risky asset risk premia targeted by habit-formation asset pricing models.

In columns 2 and 5, we include asset composition controls: the labor income/liquid wealth ratio interacted with age, the business wealth/liquid wealth ratio, and the housing wealth/liquid wealth ratio. The aim is to control for variations in background wealth. We still obtain almost identical coefficients on changes in liquid wealth. This suggests that the results are not driven by some correlation of liquid wealth changes with changes in background risk exposure due to variation in the asset mix held by the household. Our results in differences are consistent with earlier purely cross-sectional studies that have found a largely insignificant relationship between these asset composition ratios and the liquid risky asset share among stock market participants (Heaton and Lucas (2000); Yao and Zhang (2005)).

The TSLS results in columns 3 and 6 show that measurement error does not appear

⁹The simulations in Stock and Yogo (2005) suggest that with an F -statistic greater than 5.39 we can reject at a 5% significance level the hypothesis that the maximal bias of the TSLS estimator relative to OLS is greater than 0.3.

to have a major influence on our results. Both estimates are negative and close to the OLS results, in particular for the 1984 – 1999 sample, but with higher standard error. The estimate for the 1999 – 2003 subsample is somewhat larger in magnitude, but that should not be overinterpreted, because the standard error is also much larger than with OLS. The table also reports p -values from an overidentification test that show that we cannot reject that the instruments are valid (in the sense of being uncorrelated with the regression residual). Overall, the TSLS results do not provide any evidence that there is a significant positive relationship between changes in liquid wealth and changes in the liquid risky asset share.

Changes in financial risky asset shares—As an additional perspective on the issue of asset composition and background risk, Table 5 reports regressions similar to those in Table 4, but with the financial risky asset share as dependent variable, and with changes in financial wealth as explanatory variable. This perspective would be appropriate if households with CRRA preferences would keep the proportion of financial wealth invested in risky assets, including home equity and business wealth, roughly constant. In that case, the presence of habit formation would imply that changes in financial wealth should lead to changes in the financial risky asset share.

As the table shows, however, this approach doesn't produce any evidence for a positive relationship between wealth changes and risky asset shares either. The coefficients are all negative, for both subsamples, with OLS and TSLS, and with and without asset composition control (the asset composition controls here consist only of the labor income/financial wealth ratio interacted with age, as human wealth is the only remaining background wealth component that is not included in the risky asset share). The magnitudes of the coefficients are larger than in Table 4, and they are all significantly smaller than zero. Thus, the evidence is not consistent with the predictions of the habit model.

Robustness checks—Appendix A.5 reports a large number of robustness checks. The results are generally similar to those in our main tests, so we just briefly summarize here some of the variations in methodology that we explore. To check the sensitivity to the particular linearization in Eq. (8), we examine log and logit transformations of the risky asset shares. We also re-run our liquid risky asset share regressions accounting for leverage. We also find similar results when we weight observations with sample weights, and when we use a median regression estimator (which is not sensitive to outliers)

3.3 Inertia in Asset Allocation

One possible reason for the absence of a positive effect of wealth changes on risky asset shares could be that the effect is clouded by inertia. If an in- or outflow of liquid wealth materializes first in the riskless asset category (e.g., cash), and households are slow to rebalance their portfolio, this can induce a negative contemporaneous relationship

between liquid wealth changes and risky asset shares. Of course, capital gains and losses on risky assets have the opposite effect: They lead to a positive contemporaneous relationship if the household is slow to rebalance.

The regressions reported in columns 1 and 4 of Table 6 show that both effects are present. In that regression, we include a proxy for the risky asset return of the household between t and $t - k$ (for capital gains and losses, excluding dividends, to be precise), which we back out using the information on net purchases or sales of risky assets in the PSID. The risky asset return is strongly positively related to the liquid risky asset share, and changes in wealth have a more negative coefficient compared with our earlier results in Table 4, now that the risky asset return is included in the regression.

Estimating the degree of inertia—We now proceed to analyze in more detail how much inertia there is in portfolio allocations. We use the information on net purchases or sales of risky assets to construct a variable $\Delta_k \text{Inert}_t$: it represents the (counterfactual) change in the liquid risky asset share that the household would have experienced between $t - k$ and t under perfect inertia—that is, if it had not undertaken any purchases or sales of risky assets between $t - k$ and t . In this case, the risky asset position would have changed only because of capital gains and losses, and the riskless asset position would have changed only because of in- and outflows (e.g., via cash or the checking account). We then modify our wealth regression, Equation (10), by including $\Delta_k \text{Inert}_t$:

$$\Delta_k \alpha_t = \beta q_{t-k} + \gamma \Delta_k h_t + \varphi \Delta_k \text{Inert}_t + \rho \Delta_k w_t + \varepsilon_t. \quad (12)$$

If households exhibit perfect inertia, then the actual change in $\Delta_k \alpha_t$ is equal to $\Delta_k \text{Inert}_t$, and therefore $\varphi = 1$. If households exhibit no inertia at all, and hence rebalance their portfolios immediately following capital gains and inflows and outflows of liquid wealth, then $\varphi = 0$. 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 $\varphi > 1$.

It is useful to keep in mind that purchases and sales of risky assets, and hence $\Delta_k \text{Inert}_t$, are likely to be measured with significant error. In addition to the usual attenuation bias of classical errors-in-variables, the biggest concern is systematic underreporting of trades (forgotten trades). Households in the PSID are asked to recall the amount of purchases and sales over the last k years and it is plausible that they might forget some trades (Vissing-Jorgensen (2002)). In that case, part of the change in the value of liquid risky assets would be attributed wrongly to capital gains/losses instead of purchases/sales. This would lead to a spurious positive relationship between $\Delta_k \text{Inert}_t$ and $\Delta_k \alpha_t$. We do not have instruments for household-specific capital gains and losses, so using instrumental variables is not feasible and measurement error remains a concern. We can do at least a weak check by comparing results for the first subsample, where the recall period is 5 years, with those for the second subsample, where the recall period is 2 years. Also, we can look at subsamples excluding the

households that report no trades at all, which may be the most error-prone ones.

Table 6 presents results from estimating Eq. (12) with OLS. As columns 2 and 5 show, the coefficient on the inertia variable is large, around 0.75, with small standard errors. Taken at face value, it suggests that there is huge inertia. Households' asset allocations seem to fluctuate strongly as a function of in- and outflows, and capital gains and losses, without much rebalancing taking place. The coefficient on changes in liquid wealth is close to zero, as before. The R^2 is now around 0.70, which is huge compared with the small R^2 in Table 4.

But as we pointed out, it's possible that some of this effect is driven by underreporting of trades. The magnitudes of the coefficient estimates below, however, make it somewhat unlikely that this is the whole story. Underreporting would have to be extremely common to explain the magnitudes of the coefficients we find.

To provide some perspective on the trade-reporting issue, columns 3 and 6 present regressions where we interact both the inertia variable and the liquid wealth changes variable with a dummy that we name Trade_t . It takes a value of one if the household reported any net trade in risky assets for the period from $t - k$ to t and it is zero otherwise (the percentage of households that report to have traded is 59% in the 1984-1999 sample and 57% in the 1999-2003 sample). The coefficient on $\Delta_k \text{Inert}_t$ now picks up the effect for those who don't trade (and the estimate is equal to one, not surprisingly), while the effect for those who report trades can be obtained by adding the coefficient on $\Delta_k \text{Inert}_t$ and $\Delta_k \text{Inert}_t \times \text{Trade}_t$, which yields about 0.65 in both subsamples. Thus, even for those that report trades, we still find a strong inertia effect.

We also interact the Trade_t variable with changes in liquid wealth, but as the table shows there is no significant difference in the wealth changes coefficient between households that report trades and those that do not (the same is true if we exclude the inertia variable and its interaction with Trade_t from this regression). Hence, the absence of a positive effect of wealth changes on changes in risky asset shares at least is not driven by the subset of households that don't trade at all.

Allowing for slow adjustment to wealth changes—The finding that there seems to be a lot of inertia in households' portfolio shares brings up the question whether there might actually be a positive effect of wealth changes on risky asset shares, just with a time lag, because households need time to adjust, perhaps because they are trading off the benefits of rebalancing towards the optimal risky asset share against transaction costs. In Table 7 we therefore investigate the effect of wealth changes (between $t - k$ and t) on future changes in the risky asset shares (between t and $t + k$). The control variables are measured at the same points in time as earlier in Table 4, with the exception of the preference shifters and the risky asset share, which are moved k periods into the future, i.e., measured between t and $t + k$. Regarding sample requirements, we now require that no assets had been moved out of the household due to a leaving family member, and marital status remained unchanged between $t - k$ and $t + k$, and we require stock

market participation at t and $t + k$. Since we need a longer span of data for these regressions, we have a substantially lower number of observations than in Table 4.

These regressions of $\Delta_k \alpha_{t+k}$ on $\Delta_k w_t$ are also interesting from a measurement-error perspective. The concern in the earlier regressions of $\Delta_k \alpha_t$ on $\Delta_k w_t$ in Table 4 is that measurement error in riskless asset holdings (if it dominates relative to the measurement error in risky asset holdings) might induce a spurious negative relationship between w_t and α_t , and hence also between $\Delta_k \alpha_t$ and $\Delta_k w_t$. However, for the regressions in Table 7 the situation is different: If measurement error induces a mechanical negative relationship between w_t and α_t , it should lead to a spurious *positive* relationship between $\Delta_k \alpha_{t+k} \equiv \alpha_{t+k} - \alpha_t$ and $\Delta_k w_t \equiv w_t - w_{t-k}$. This is easiest to see with when measurement error is assumed to be uncorrelated over time, but it is also true with positively autocorrelated measurement error.

As Table 7 shows, the point estimates for the wealth effect are indeed positive, and statistically significant in the first subsample (1984 – 1999) but not in the second (1999 – 2003). However, in terms of economic magnitudes, the coefficient estimates are again close to zero and not much different from those in Table 4. If we take the maximum coefficient estimate (0.040, column 1) in the table, it suggests that an increase in liquid wealth by 10% leads to an increase in the risky asset share from 50% to 50.4%, which strikes us as a small effect. Moreover, if one was concerned that the estimates in Table 4 might have a negative measurement error bias, then the estimates in Table 7 would have positive measurement error bias and would therefore overstate the effect of wealth changes.¹⁰ Overall, these results suggest that even if we allow for slow adjustment, there is no evidence for an economically significant effect of liquid wealth changes on risky asset shares.

Big vs. small changes—One possible explanation for inertia is that households face some fixed rebalancing cost. In that case, households would only want to rebalance if the benefits are large enough to outweigh the fixed rebalancing cost. From the perspective of the habit-formation model, this would imply that the household might be unwilling to rebalance following small wealth changes, but it might do so after big changes.

To find out, we examine piecewise-linear regressions, shown in Figure 1. We run regressions similar to those in Table 4, with the full set of controls (except the asset-composition controls), and we use a spline for $\Delta_k w_t$. We set spline breakpoints at the quartiles of the distribution of $\Delta_k w_t$. Panel (a) in Figure 1 presents the fitted values, where we express $\Delta_k w_t$ relative to its median and normalize such that the lines cross the origin. The range of values shown for $\Delta_k w_t$ in the graph is about twice the difference between the 75th and 25th percentile. We omit standard errors from the graph, but

¹⁰We also estimated the regressions in Table 7 with TSLS. The estimates are close to the OLS estimates. But due to the lower number of observations, the instruments are now very weak in the first stage so that the second-stage estimates are not reliable. For this reason, we do not report the TSLS results.

none of the slopes in the four segments is more than two standard errors from zero. As the figure shows, in both the 1984-1999 and the 1999-2003 sample, the relationship between the liquid risky asset share and changes in liquid wealth is flat for small and large values of $\Delta_k w_t$. Hence, there is no support for the view that households might conform more closely to the predictions of the habit model when wealth changes are big and hence the benefits from rebalancing towards the optimal portfolio should be large.

More generally, one can also ask whether households might exhibit less inertia after big in-/outflows or big capital gains/losses. Again, with fixed rebalancing costs households might be reluctant to rebalance unless the asset allocation has moved sufficiently far away from the optimum. Therefore, Panel (b) in Figure 1 presents a piecewise-linear version of the regressions in column 1 and 3 of Table 6. Spline breakpoints are now set at the quartiles of the distribution of $\Delta_k \text{Inert}_t$. Everything else, too, is similar to Panel (a), just with $\Delta_k \text{Inert}_t$ replacing $\Delta_k w_t$. Standard errors are again omitted to reduce clutter, but they are small relative to the point estimates of the slope coefficients (between 0.06 and 0.13 for the slope coefficients in each of the four segments). The figure shows that inertia is weaker when $\Delta_k \text{Inert}_t$ is above the 75th percentile, i.e., after big capital gains or following large outflows. However, inertia is still relatively strong, and statistically still clearly different from zero.

That there is inertia even after big in-/outflows or big capital gains/losses casts some doubt on the explanation that households are trading off fixed rebalancing costs against benefits of rebalancing. It seems more likely that households simply do not pay close attention to their portfolio allocations. In other words, the costs of devoting any attention to the portfolio may be important, rather than actual costs of transacting.

4 Discussion

Summing up, our evidence shows that the effect of wealth changes on households' asset allocation predicted by difference-habit models is absent in microdata. The relationship between wealth and asset allocation seems best described by constant relative risk aversion. However, the large inertia we find isn't predicted by constant relative risk aversion models either—at least not without adding frictions.

Our focus in this paper is on understanding the microeconomics of household asset allocation. But beyond this microeconomic perspective, our results also raise some questions about models with habit-formation in asset pricing and macroeconomics. In difference-habit asset-pricing models, variations in aggregate wealth over the business cycle generate large low-frequency variation in relative risk aversion and the relative demands for risky and riskless assets. However, the household-specific variation in wealth that we see in microdata seems rather large relative to business-cycle variation, and should therefore generate even larger household-specific variation in relative

risk aversion and asset allocation. We don't find this variation in asset allocation in microdata.

To be clear, we cannot directly test the microeconomic implications of representative agent models like Campbell and Cochrane (1999) because it is not even clear how the microfoundations of these models would look like, except for some special cases with complete markets. However, notwithstanding this lack of explicit microfoundations, researchers often view the preferences of the representative agent in these models as being a plausible representation of the preferences of microeconomic agents. For example, Campbell and Cochrane (1999) motivate their choice of habit-formation preferences by pointing out that they are appealing from a psychological perspective, and, in particular, that the microeconomic predictions of external habits for consumption are plausible. Our findings with microdata cast doubt on the plausibility of such microeconomic stories for time-varying risk aversion.

Our finding that wealth changes have some impact on stock market entry and exit suggests that changing stock market participation, rather than time-varying individual risk aversion, could perhaps play a role in the time-variation of risk premia in the aggregate. Wealth changes can induce changes in stock market participation even with CRRA preferences, if there are some per-period participation costs. But the effect we find does not seem very strong, so it is somewhat questionable whether the magnitudes are big enough to have a significant effect in the aggregate.

Finally, the strong asset allocation inertia we find is an interesting and, so far, not well-understood phenomenon.¹¹ At a given point in time, a household's asset allocation depends to a large extent on the past history of capital gains/losses and in-/outflows. Part of it may reflect underreporting of risky asset purchases and sales in the PSID, but we doubt that such measurement error can explain the bulk of the apparent inertia, not least because similar inertia has also been found with data from 401(k) retirement accounts that do not have the same measurement error problems. Samuelson and Zeckhauser (1988), Ameriks and Zeldes (2001), Agnew, Balduzzi, and Sunden (2004), 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.

One explanation could be that individuals are not willing to rebalance their portfolios because they perceive it as too costly. If so, it seems to be more a cost of giving any attention at all to the portfolio, rather than a fixed rebalancing cost in the form of explicit transaction costs, because we find that households are almost as reluctant to rebalance following large wealth changes as they are after small wealth changes.

In any case, slow adjustment of portfolio shares does not explain the absence of a

¹¹A recent paper by Biliias, Georgarakos, and Haliassos (2006) explores household characteristics that are correlated with portfolio inertia. Note that our finding of asset allocation inertia at the portfolio level is not in contradiction with findings in Odean (1998) at the individual stock level that investors tend to sell stocks with good past performance (the disposition effect).

wealth effect on risky asset shares in our data, because wealth changes do not have an economically significant effect on future changes in risky asset shares either.

References

- ABEL, A. B. (1990): “Asset Prices under Habit Formation and Catching Up with the Joneses,” *American Economic Review*, 80(2), 38–42.
- AGNEW, J., P. BALDUZZI, AND A. SUNDEN (2004): “Portfolio Choice and Returns in a Large 401(k) Plan,” *American Economic Review*, forthcoming.
- AMERIKS, J., AND S. P. ZELDES (2001): “How Do Household Portfolio Shares Vary With Age?,” working paper, Columbia University.
- BAKSHI, G. S., AND Z. CHEN (1996): “The Spirit of Capitalism and Stock-Market Prices,” *American Economic Review*, 86(1), 133–157.
- BARKSY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, 112(2), 537–579.
- BERTAUT, C. C., AND M. HALIASSOS (1995): “Why Do so Few Hold Stocks?,” *Economic Journal*, 105(432), 1110–1129.
- BILIAS, Y., D. GEORGARAKOS, AND M. HALIASSOS (2006): “Portfolio Inertia and Stock Market Fluctuations,” working paper, Goethe University Frankfurt.
- BOLDRIN, M., L. J. CHRISTIANO, AND J. D. FISHER (2001): “Habit Persistence, Asset Returns, and the Business Cycle,” *American Economic Review*, 91(1), 149–166.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2), 205–251.
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), 1–55.
- CHAMBERLAIN, G. (1984): “Panel Data,” in *Handbook of Econometrics*, ed. by Z. Griliches, and M. D. Intriligator, vol. 2, pp. 1247–1318. North Holland, Amsterdam.

- CHETTY, R., AND A. SZEIDL (2005): “Consumption Commitments: Neoclassical Foundations for Habit Formation,” working paper, UC Berkeley.
- COCCO, J. F. (2004): “Portfolio Choice in the Presence of Housing,” *Review of Financial Studies*, 18(2), 535–567.
- COCCO, J. F., P. MAENHOUT, AND F. GOMES (2005): “Consumption and Portfolio Choice over the Life-Cycle,” *Review of Financial Studies*, 18(2), 491–533.
- CONSTANTINIDES, G. M. (1990): “Habit Formation: A Resolution of the Equity Premium Puzzle,” *Journal of Political Economy*, 98(3), 519–543.
- DEATON, A. (1992): *Understanding Consumption*. Clarendon Press, Oxford.
- (1997): *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Johns Hopkins University Press, Baltimore, MD.
- DYNAN, K. E. (2000): “Habit Formation in Consumer Preferences: Evidence from Panel Data,” *American Economic Review*, 90(3), 391–406.
- FAIG, M., AND P. SHUN (2002): “Portfolio Choice in the Presence of Personal Illiquid Projects,” *Journal of Finance*, 57(1), 303–328.
- GOLLIER, C., AND J. W. PRATT (1996): “Risk Vulnerability and the Tempering Effect of Background Risk,” *Econometrica*, 64(5), 1109–1123.
- GOMES, F., AND A. MICHAELIDES (2003): “Portfolio Choice with Internal Habit Formation: A Life-Cycle Model with Uninsurable Labor Income Risk,” *Review of Economic Dynamics*, 6(4), 729–766.
- (2005): “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *Journal of Finance*, 60(2), 869–904.
- GUISSO, L., M. HALIASSOS, AND T. JAPPELLI (2003): “Household Stockholding in Europe: Where Do We Stand and Where Do We Go?,” *Economic Policy*, 18(37), 523–577.
- HEATON, J., AND D. LUCAS (1997): “Market Frictions, Savings Behavior, and Portfolio choice,” *Macroeconomic Dynamics*, 1(1), 76–101.
- (2000): “Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk,” *Journal of Finance*, 55(3), 1163–1198.
- HUBERMAN, G., AND P. SENGMUELLER (2004): “Performance Predicts Asset Allocation: Company Stock in 401(k) Plans,” *Review of Finance*, 8(3), 403–443.

- JERMANN, U. (1998): “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, 41(2), 257–275.
- JUSTER, F. T., J. P. SMITH, AND F. STAFFORD (1999): “The Measurement and Structure of Household Wealth,” *Labour Economics*, 6, 253–275.
- LUPTON, J. P. (2003): “Household Portfolio Choice and the Habit Liability: Evidence from Panel Data,” working paper, Federal Reserve Board.
- MANKIW, N. G., AND S. ZELDES (1991): “The Consumption of Stockholders and Nonstockholders,” *Journal of Financial Economics*, 29(1), 97–112.
- MEER, J., D. L. MILLER, AND H. S. ROSEN (2003): “Exploring the Health-Wealth Nexus,” working paper, National Bureau of Economic Research.
- MEHRA, R., AND E. C. PRESCOTT (2003): “The Equity Premium in Retrospect,” in *Handbook of the Economics of Finance*, ed. by G. M. Constantinides, M. Harris, and R. Stulz, pp. 887–936. Elsevier North-Holland, Amsterdam.
- MOREIRA, M. (2003): “A Conditional Likelihood Ratio Test for Structural Models,” *Econometrica*, 71(4), 1027–1048.
- ODEAN, T. (1998): “Are investors reluctant to realize their losses?,” *Journal of Finance*, 53(5), 1887–1934.
- RAVINA, E. (2005): “Keeping Up with the Joneses: Evidence from Micro Data,” working paper, Northwestern University.
- SAHM, C. R. (2006): “Does Risk Tolerance Change?,” working paper, University of Michigan.
- SAMUELSON, P. A. (1969): “Lifetime Portfolio Selection By Dynamic Stochastic Programming,” *Review of Economics and Statistics*, 51(3), 239–246.
- SAMUELSON, W., AND R. ZECKHAUSER (1988): “Status Quo Bias in Decision Making,” *Journal of Risk and Uncertainty*, 1(1), 7–59.
- STAIGER, D., AND J. H. STOCK (1997): “Instrumental Variables Regression with Weak Instruments,” *Econometrica*, 65, 557–586.
- STOCK, J. H., AND M. YOGO (2005): “Testing for Weak Instruments in Linear IV Regression,” in *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, ed. by D. W. Andrews, and J. H. Stock, pp. 80–108. Cambridge University Press, Cambridge.

VISSING-JØRGENSEN, A. (2002): “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structure,” working paper, Northwestern University.

YAO, R., AND H. H. ZHANG (2005): “Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints,” *Review of Financial Studies*, 18(1), 197–239.

ZELDES, S. P. (1989): “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *Quarterly Journal of Economics*, 104(2), 275–298.

Table 1: Summary Statistics

Variable	Mean	10 th pct.	Median	90 th pct.	N
All Households, 1984 - 1999 ($k = 5$ years)					
Liquid wealth	156,391	480	52,532	349,116	3,313
Financial wealth	430,238	26,728	184,756	830,154	3,313
Income	92,515	24,367	72,760	160,369	3,319
Stock mkt. particip.	0.45	0	0	1	3,319
Stock mkt. entry	0.35	0	0	1	1,408
Stock mkt. exit	0.24	0	0	1	1,909
All Households, 1999 - 2003 ($k = 2$ years)					
Liquid wealth	201,783	1,003	62,936	399,129	3,035
Financial wealth	479,317	24,989	202,029	890,536	3,035
Income	99,665	26,884	76,397	172,540	3,033
Stock mkt. particip.	0.58	0	1	1	3,035
Stock mkt. entry	0.34	0	0	1	888
Stock mkt. exit	0.19	0	0	1	2,147
Stock Market Participants, 1984 - 1999 ($k = 5$ years)					
Liquid wealth	269,609	19,137	101,827	576,663	1,439
Financial wealth	630,488	71,442	286,508	1,155,371	1,439
Income	118,502	37,917	90,570	196,475	1,439
Δ_k log liq. wealth	0.36	-0.98	0.47	1.62	1,399
Δ_k log fin. wealth	0.30	-0.53	0.29	1.16	1,429
Δ_k log income	0.05	-0.52	0.09	0.60	1,439
%liq. assets risky	0.56	0.13	0.57	0.95	1,439
%fin. wealth risky	0.74	0.36	0.78	0.99	1,438
Δ_k %liq. assets risky	0.09	-0.37	0.06	0.59	1,439
Δ_k %fin. wealth risky	0.05	-0.30	0.03	0.41	1,438
Stock Market Participants, 1999 - 2003 ($k = 2$ years)					
Liquid wealth	294,622	16,904	110,980	550,204	1,710
Financial wealth	640,382	63,540	296,664	556,920	1,710
Income	116,432	34,000	90,126	198,552	1,710
Δ_k log liq. wealth	-0.04	-1.29	-0.02	1.24	1,654
Δ_k log fin. wealth	0.09	-0.74	0.10	0.88	1,694
Δ_k log income	-0.09	-0.62	0.00	0.39	1,710
%liq. assets risky	0.58	0.17	0.59	0.96	1,710
%fin. wealth risky	0.75	0.39	0.81	0.99	1,710
Δ_k %liq. assets risky	-0.02	-0.46	-0.01	0.40	1,710
Δ_k %fin. wealth risky	0.00	-0.32	0.00	0.33	1,710

Table 2: Changes in Liquid Wealth and Stock Market Entry and Exit: Probit Regressions

	Entry		Exit	
	$k = 5$	$k = 2$	$k = 5$	$k = 2$
	(1984-1999)	(1999-2003)	(1984-1999)	(1999-2003)
$\Delta_k \log \text{liq. wealth}_t$	0.124 (0.014)	0.108 (0.016)	-0.058 (0.006)	-0.072 (0.007)
$\Delta_2 \log \text{income}_{t-k}$	0.048 (0.032)	0.019 (0.035)	-0.051 (0.021)	-0.034 (0.015)
$\Delta_2 \log \text{income}_{t-k-2}$	0.009 (0.024)	0.004 (0.026)	-0.024 (0.018)	-0.020 (0.011)
$\log \text{income}_{t-k-4}$	0.203 (0.039)	0.122 (0.040)	-0.097 (0.023)	-0.044 (0.014)
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
Pseudo R^2	0.22	0.18	0.15	0.19
N	971	607	1,556	1,724

Notes: Estimates are marginal effects evaluated at sample averages of the explanatory variables. Standard errors are reported in parentheses.

Table 3: First Stage Regressions

	$\Delta_k \log \text{liq. wealth}_t$		$\Delta_k \log \text{fin. wealth}_t$	
	$k = 5$ (1984-1999)	$k = 2$ (1999-2003)	$k = 5$ (1984-1999)	$k = 2$ (1999-2003)
Instruments:				
$I_{(\Delta_k \log \text{income}_t < 10^{th} \text{ pct.})}$	-0.211 (0.131)	-0.110 (0.093)	-0.217 (0.093)	-0.116 (0.073)
$I_{(\Delta_k \log \text{income}_t > 90^{th} \text{ pct.})}$	0.480 (0.121)	0.098 (0.082)	0.257 (0.082)	0.062 (0.067)
Inheritance $_t$	0.290 (0.160)	0.559 (0.174)	0.111 (0.082)	0.417 (0.104)
Controls:				
$\Delta_2 \log \text{income}_{t-k}$	0.165 (0.101)	0.087 (0.057)	0.077 (0.048)	0.033 (0.046)
$\Delta_2 \log \text{income}_{t-k-2}$	0.141 (0.120)	0.092 (0.043)	0.085 (0.043)	0.033 (0.029)
Log income $_{t-k-4}$	0.133 (0.102)	0.068 (0.042)	0.040 (0.045)	0.014 (0.031)
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
Partial R^2 of instruments	0.01	0.01	0.02	0.01
F-test of instruments	8.39	5.09	6.19	7.41
[p-value]	[0.00]	[0.00]	[0.00]	[0.00]
N	1,234	1,455	1,258	1,489

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses.

Table 4: Changes in the Proportion of Liquid Assets Invested in Risky Assets

	$k = 5$ (1984-1999)			$k = 2$ (1999-2003)		
	OLS	OLS	TOLS	OLS	OLS	TOLS
$\Delta_k \log \text{liquid wealth}_t$	-0.013 (0.007)	-0.009 (0.009)	-0.012 (0.058)	0.023 (0.011)	0.017 (0.015)	-0.136 (0.076)
Asset composition controls		Y			Y	
Preference shifters	Y	Y	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y	Y	Y
Adj. R^2	0.01	0.01	–	0.01	0.02	–
Overidentification test	–	–	[0.41]	–	–	[0.64]
N	1,234	1,234	1,234	1,455	1,455	1,455

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses, p -values in brackets.

Table 5: Changes in the Proportion of Financial Wealth Invested in Risky Assets

	$k = 5$			$k = 2$		
	(1984-1999)			(1999-2003)		
	OLS	OLS	TSLS	OLS	OLS	TSLS
$\Delta_k \log \text{financial wealth}_t$	-0.160 (0.059)	-0.172 (0.091)	-0.198 (0.090)	-0.108 (0.031)	-0.103 (0.036)	-0.355 (0.130)
Asset composition controls		Y			Y	
Preference shifters	Y	Y	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y	Y	Y
Adj. R^2	0.11	0.11	–	0.06	0.06	–
Overidentification test	–	–	[0.56]	–	–	[0.57]
N	1,258	1,258	1,258	1,489	1,489	1,489

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses, p -values in brackets.

Table 6: Effects of Inertia on Changes in the Proportion of Liquid Assets Invested in Risky Assets, OLS

	$k = 5$ (1984-1999)			$k = 2$ (1999-2003)		
$\Delta_k \log \text{liq. wealth}_t$	-0.061 (0.025)	0.000 (0.005)	-0.003 (0.003)	-0.167 (0.014)	0.005 (0.006)	-0.001 (0.002)
$\Delta_k \log \text{liq. wealth}_t \times \text{Trade}_t$			0.001 (0.002)			0.003 (0.010)
Risky asset return $_t$	0.151 (0.012)			0.227 (0.013)		
$\Delta_k \text{Inert}_t$		0.743 (0.027)	1.002 (0.010)		0.754 (0.054)	1.004 (0.006)
$\Delta_k \text{Inert}_t \times \text{Trade}_t$			-0.347 (0.037)			-0.369 (0.068)
Trade $_t$			0.128 (0.011)			0.021 (0.010)
Preference shifters	Y	Y	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y	Y	Y
Adj. R^2	0.34	0.64	0.70	0.57	0.72	0.76
N	1,042	1,080	1,080	1,308	1,325	1,325

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses.

Table 7: Future Changes in the Proportion of Liquid Assets Invested in Risky Assets:
 $\Delta_k \alpha_{t+k}$ as Dependent Variable, OLS

	$k = 5$ (1984-1999)		$k = 2$ (1999-2003)	
$\Delta_k \log \text{liquid wealth}_t$	0.040	0.037	0.006	0.013
	(0.015)	(0.015)	(0.015)	(0.014)
Asset composition controls		Y		Y
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
Adj. R^2	0.00	0.00	0.00	0.02
N	561	561	597	597

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses.

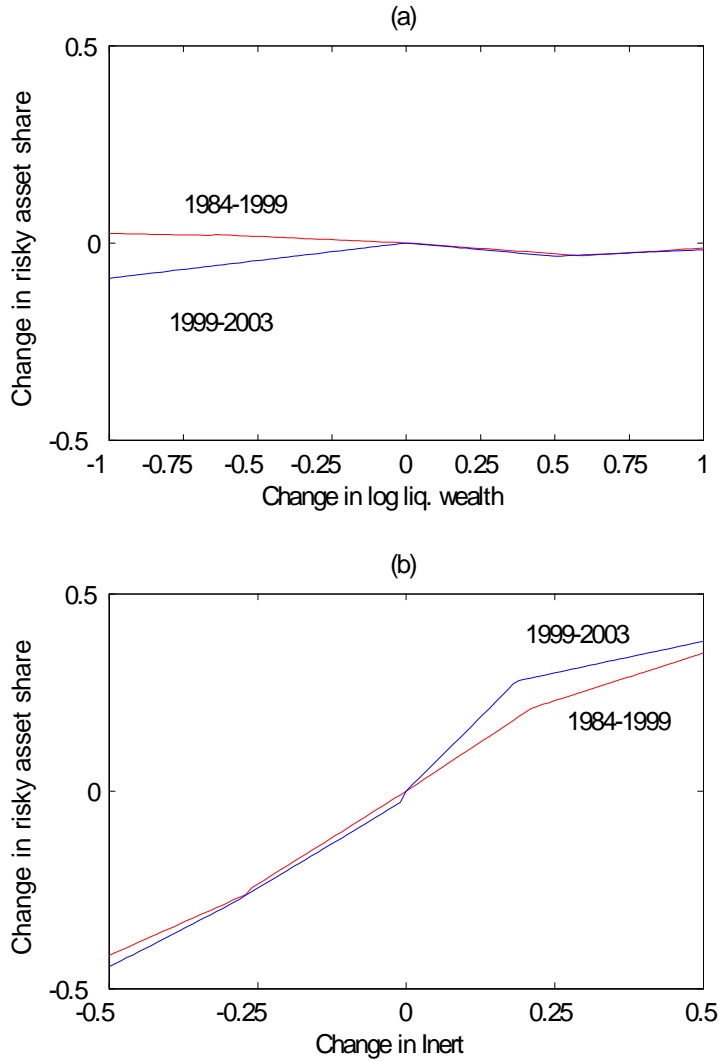


Figure 1: Piecewise-Linear Regression of Changes in the Liquid Risky Asset Share on Changes in Liquid Wealth (a), and on the Inertia Variable (b).

Notes: In addition to the spline terms for $\Delta_k w_t$ or $\Delta_k \text{Inert}_t$, the regressions also contain the full set of control variables (Life-cycle controls, preference shifters, and year-region FE, as in Table 4). Spline breakpoints are set at quartiles for $\Delta_k w_t$ or $\Delta_k \text{Inert}_t$, respectively. For the plot, we express $\Delta_k w_t$ and $\Delta_k \text{Inert}_t$ relative to their medians. Control variables are held fixed as we vary $\Delta_k w_t$ and $\Delta_k \text{Inert}_t$. The figure presents fitted values normalized such that the lines cross the origin.

APPENDIX

A.1 Model with Internal Habit

The analysis in Section 2.1 can easily be extended to models with internal habit, where the habit depends on past consumption as in Constantinides (1990) in a continuous-time setting. Specifically, let the habit, X_t , follow the difference equation:

$$X_{t+1} - X_t = bC_t - aX_t. \quad (\text{A.1})$$

Define

$$C_t^* \equiv (C_t - X_t) \frac{R_f + a}{R_f - b + a}, \quad (\text{A.2})$$

and

$$W_t^* \equiv W_t - (1 + R_f) \frac{X_t}{R_f - b + a}. \quad (\text{A.3})$$

W_t^* reflects the excess wealth that is not needed to finance future discounted habit, where the habit grows at a rate of b and depreciates at a rate of a . Note that the value of X_{t+1} is known at time t . Now assume that the investor at time t invests, after time t consumption, a fraction α_t^* of wealth in excess of $\frac{X_{t+1}}{R_f - b + a}$ into the risky asset, and the rest in the riskless asset. This surplus portfolio yields a return $R_{p,t+1}^* \equiv \alpha_t^* (R_t - R_f) + R_f$. The remaining $\frac{X_{t+1}}{R_f - b + a}$ dollars are invested in the riskless asset. The dynamic budget constraint becomes

$$W_{t+1} = (1 + R_{p,t+1}^*) \left(W_t - C_t - \frac{X_{t+1}}{R_f - b + a} \right) + (1 + R_f) \frac{X_{t+1}}{R_f - b + a}. \quad (\text{A.4})$$

Substituting in the definitions of surplus wealth and consumption and multiplying terms we get

$$W_{t+1}^* = (1 + R_{p,t+1}^*) (W_t^* - C_t^*). \quad (\text{A.5})$$

Using the definition of C_t^* the optimization problem now is

$$\max E_t \sum_{\tau=0}^{\infty} \delta^\tau \frac{C_{t+\tau}^{*1-\gamma}}{1-\gamma} \left(\frac{R_f - b + a}{R_f + a} \right)^{1-\gamma}, \quad (\text{A.6})$$

which is equivalent to

$$\max E_t \sum_{\tau=0}^{\infty} \delta^\tau \frac{C_{t+\tau}^*}{1-\gamma}, \quad (\text{A.7})$$

so the problem again maps into a power utility problem. The portfolio share now is

$$\alpha_t = \alpha^* \left(1 - \frac{X_{t+1}}{(W_t - C_t)(R_f + a - b)} \right). \quad (\text{A.8})$$

Approximating $\alpha^* \approx 1$, as in the main text, we get

$$\alpha_t = 1 - \frac{X_{t+1}}{(W_t - C_t)(R_f + a - b)}. \quad (\text{A.9})$$

Log-linearizing and, in an abuse of notation, slightly changing our definitions to $x_t \equiv \log(X_t/(R_f + a - b))$, $w_t \equiv \log(W_t - C_t)$, we have

$$\alpha_t = 1 - \exp(x_{t+1} - w_t) \quad (\text{A.10})$$

$$\approx \kappa - \rho(x_{t+1} - w_t), \quad (\text{A.11})$$

which implies

$$\Delta\alpha_t = -\rho\Delta x_{t+1} + \rho\Delta w_t. \quad (\text{A.12})$$

Note that $\Delta x_{t+1} = \Delta \log(X_{t+1}) = \log(1 + bC_t - aX_t)$. Therefore, as long as b and a are close to zero, which means that the habit reacts sluggishly to past consumption, we can approximate

$$\Delta\alpha_t \approx \rho\Delta w_t. \quad (\text{A.13})$$

as we do in Section 2.2. Of course, if the habit reacts faster, then it is possible that the effects of Δx_{t+1} and Δw_t offset, and our tests don't pick up the time-varying risk aversion induced by the habit. So our tests should be viewed as tests for low frequency movements in relative risk aversion.

A.2 Numerical Solution of a Model with Anticipated Income

To illustrate the effects of anticipated income on portfolio choice, we numerically solve a three-period version of our model, in which the household receives, with some positive probability, a large payment (e.g., an inheritance) in the second period. The household starts with an initial wealth of $W_0 = 50$ in liquid assets and chooses consumption (C_1, C_2) and the risky asset share (α_1, α_2) at time $t = 1, 2$. For simplicity, the household is assumed to have no labor income or other assets. At $t = 3$, the household is assumed to consume the entire remaining wealth. At $t = 2$, the household receives a payment B , with probability $p_B = 0.8$. Hence, the expected value of this payment to the household at $t = 1$ is $E_1[B] = p_B B$. We further set $\log(1 + R_t) \sim \mathcal{N}(\mu, \sigma^2)$, $\log(1 + R_f) = 0.04$, $\mu = 0.09$, $\sigma = 0.15$, $\gamma = 4$, and $\delta = 0.9$.

We solve the model by backward induction using a standard approach. We first solve the second period problem as a function of beginning of second period wealth. For a given level of beginning of second period wealth (before second period consumption but after returns from the first to the second period are realized), we perform a grid search over values of C_2 and α_2 to find the combination that maximizes expected utility,

where we use numerical integration to evaluate expected utility. In this way we obtain maximized expected utility, i.e., the value function, as a function of beginning of second period wealth on a discretely spaced grid. We interpolate the value function between grid points and then solve the first period problem to obtain the optimal C_1 and α_1 .

Figure A.1 shows how consumption and the risky asset share chosen in the first period depend on the expected value of the second period payment. We consider values for B from 0 to 50, i.e., $E_1[B]$ ranges from 0 to 40. At the higher end of this range, the payment, if received, substantially raises the liquid assets of the household (compared with $W_0 = 50$), and this increase is largely anticipated since $p_B = 0.8$ —just like it might be the case for the typical inheritance. Panel (a) shows that the risky asset share of a household with habit ($X = 10$) does not significantly increase as we increase B . Despite the fact that the household anticipates a substantial asset inflow in the second period, this does not induce the household to increase the allocation to risky asset. The reason is—as we discuss in the main part of the paper—that the small, but non-negligible risk that B will be zero forces the household to still save enough in riskless assets to be able to self-insure future habit. A likely, but not entirely certain payment in the second period cannot be utilized for insuring future habit, and therefore does not significantly increase the household’s willingness to hold the risky asset. In fact, there is actually a small decline in the risky asset share with higher B . This effect has to do with the household’s consumption decision. As Panel (b) shows, and as one would expect, the anticipation of a large payment raises consumption in the first period, which in turn implies that less liquid assets are available to insure future habit, and so a higher proportion of those liquid assets must be invested in the riskless asset.

For comparison, Panels (a) and (b) also show the optimal consumption and risky asset share of a household without habit ($X = 0$). As one would expect, the promise of a large payment raises the willingness of a CRRA household to hold risky assets in the first period, because the risky inflow in the second period provides some diversification of the risk associated with the risky asset, akin to the effect of risky, uncorrelated labor income.

Panels (c) and (d) plot $\Delta\alpha_2 \equiv \alpha_2 - \alpha_1$ against the change in the log of liquid assets $\Delta w_2 = w_2 - w_1$, comparable to the variables that we measure in the empirical data, for several values of B . It is apparent that the risky asset share of a household with habit utility strongly responds to changes in the level of liquid assets (Panel (c)), while the risky asset share of a household with CRRA utility does not (Panel (d)). Most importantly, with habit utility the relationship is almost identical, irrespective of whether changes in the level of liquid assets are unexpected ($E_1[B] = 0$) or a large increase is anticipated ($B = 50$, $E_1[B] = 40$). Overall, the results from this model support our intuitive argument in Section 2.2 of the paper that the effects of unexpected and anticipated changes in liquid asset holdings on the risky asset share should be similar.

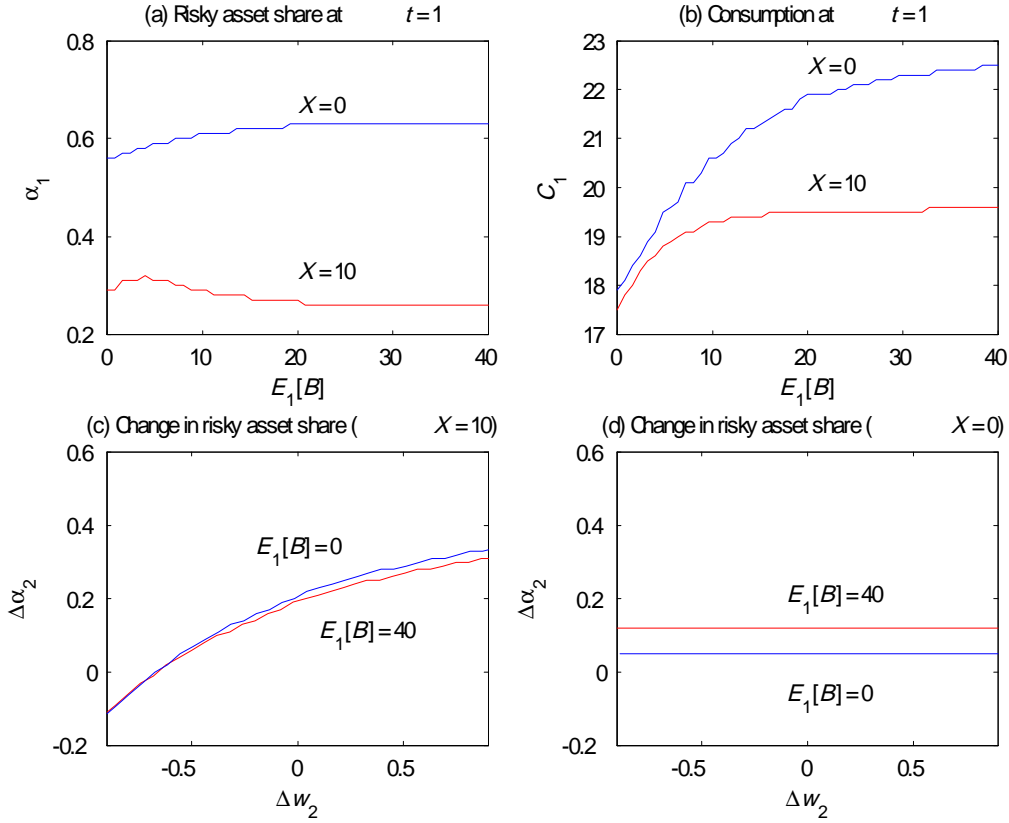


Figure A.1: Numerical Solutions of Three-Period Model

A.3 Estimation in the Presence of Aggregate Shocks

Chamberlain (1984) points out that when households are subject to common aggregate shocks, the fact that a model implies that the time-series average of shocks converges to zero as $T \rightarrow \infty$ (a typical implication of rational expectations models), does not imply that the cross-sectional average must go to zero as the number of cross-sectional units $N \rightarrow \infty$. If the shock is identical across the population, then including time dummies absorbs the common shock. But if different groups of households have different sensitivity to the common shock, there is again no guarantee that the time dummies eliminate the common shock. Hence, estimation with cross-sectional moments may lead to biased coefficients. In our setting, however, this issue does not arise. To show this, we closely follow Deaton (1992), p. 148. Assume that household i 's log wealth is subject to an aggregate shock ξ_t with household-specific sensitivity $(1 + \eta_i)$ and an idiosyncratic shock (uncorrelated across households) ν_{it} , so that

$$\Delta w_{it} = \xi_t + \eta_i \xi_t + \nu_{it}, \quad (\text{A.14})$$

where the cross-sectional average of η_i is zero. Our model implies that the change in the risky asset share follows

$$\Delta \alpha_t = \rho(\eta_i \xi_t + \nu_{it}) + \lambda \rho \xi_t, \quad 0 < \lambda < 1, \quad (\text{A.15})$$

where the coefficient on the aggregate shock is lower by the factor λ , because ξ_t implies a change in the aggregate demand for stocks which leads to a change in prices and expected returns, not to a change in the quantity of stocks held (holding the supply of stocks fixed). Now suppose we run a single cross-sectional regression, with intercept (i.e., a time dummy), of $\Delta \alpha_t$ on Δw_{it} . It is easy to show that the OLS estimator for the slope coefficient has the probability limit

$$\text{plim}_{N \rightarrow \infty} \hat{\beta} = \frac{\rho \text{Var}(\eta_i \xi_t + \nu_{it})}{\text{Var}(\eta_i \xi_t + \nu_{it})} = \rho, \quad (\text{A.16})$$

where $\text{Var}(\cdot)$ denotes the cross-sectional variance. Thus, we can consistently estimate ρ by including time-dummies in our panel regressions.

The reason that the Chamberlain (1984) problem does not arise in our setting is that dependent and explanatory variable are contemporaneous. As a consequence, the effects of $\eta_i \xi_t$ in the dependent and explanatory variable cancel out in numerator and denominator of $\text{plim} \hat{\beta}$. In contrast, the examples discussed by Deaton (1992), p. 146-148, are ones where consumption growth is regressed on *lagged* income growth.

One issue that may complicate things is if λ is also heterogeneous in the population and is correlated with η_i . Then it no longer drops out completely through the time dummy. With $\Delta \alpha_t$ as the LHS variable, λ should have some heterogeneity, because a given percentage change in stock prices due to the aggregate shock should have a bigger effect for households with high α_t than for those with low α_t . However, we can also linearize Eq. (8) differently so that $\Delta \log \alpha_t$ is on the LHS. In that case, the effect of a given percentage change in prices on $\Delta \log \alpha_t$ would be the same for all households (assuming the composition of their risky asset portfolios is similar). We report regressions with $\Delta \log \alpha_t$ in Table A.2 in this Appendix. They yield similar results compared with those that have $\Delta \alpha_t$ as the dependent variable, which suggests that heterogeneity in λ does not have a significant effect on our results.

A.4 Data: Panel Study of Income Dynamics

We now provide additional details about our data from the PSID and variable construction. Whenever possible, we use the Wealth Supplement Files and the Income Plus Files to construct our variables, and the Core Family Files otherwise. Annual

sample sizes in the PSID range from 5,000 to 7,000, but they are significantly reduced by the data availability requirements we impose.

In terms of timing, wealth data is reported as of the time the interview takes place (e.g., some time during 2003 in the 2003 wave), while income data refer to the calendar year preceding the date of the interview. Hence, the income and wealth data are not perfectly aligned, but for our tests this does not constitute a problem, because we focus on the relationship of different wealth variables which are all measured at the same date for a given household.

Riskless assets comprise the PSID categories cash (checking and savings accounts, money market funds, certificates of deposits, savings bonds, and treasury bills) plus bonds and life insurance (bonds, bond funds, cash value in a life insurance, valuable collection for investment purposes, and rights in a trust or estate). Risky liquid assets are defined as the amount reported in the PSID survey question asking for the combined value of shares of stock in publicly held corporations, mutual funds, and investment trusts. In the PSID, subjects are asked to report securities holdings net of amounts owed on the position. Other debts comprise items such as credit card debt, student loans, medical or legal bills, and loans from relatives. Home equity is the value of the home minus remaining mortgage principal.

Before 1999, subjects were asked explicitly to include assets held in individual retirement accounts (IRA) when reporting their financial asset holdings. Since 1999, they are 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.

When we use data on purchases and sales of risky assets to back out capital gains and losses, 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. For IRA assets in 1999, 2001, and 2003 we only have a combined active investment figure for all IRA assets. As an approximation, we assume that new IRA funds are allocated pro rata among the prior holdings.

A.5 Robustness Checks

Correction for weak instruments—In our TSLS regressions, our instruments are highly significant in the first-stage regression, but comparing the first stage F -statistic for the test that the coefficients on the instruments are jointly zero with the results of Stock and Yogo (2005) nevertheless suggests that we cannot reject with high confidence that the

TOLS estimator could have some bias and the test statistics could have size distortions. For this reason, we re-run our tests with a limited information maximum likelihood (LIML) estimator and compute coverage-corrected confidence intervals along the lines proposed by Moreira (2003). Table (A.1) presents the results. The point estimate shown is the LIML estimate and the coverage-corrected 95% confidence interval is shown in brackets. In Specification (1) the dependent variable is the change in the proportion of risky assets in liquid wealth, and in Specification (2) the dependent variable is the change in the proportion of risky assets in financial wealth. Comparing the results with Tables 4 and 5, it is apparent that the LIML estimator produces results that are almost identical to those with the TOLS estimator. The only difference is that the coverage-corrected confidence intervals are slightly wider than the TOLS confidence intervals based on the usual normal approximation. Overall, our TOLS results do not seem to be much affected by a weak-instruments problem.

Table A.1: LIML estimates and coverage-corrected Moreira (2003) confidence intervals

	$k = 5$		$k = 2$	
	(1984-1999)		(1999-2003)	
	(1)	(2)	(1)	(2)
$\Delta_k \log \text{ liquid wealth}_t$	-0.012 [-0.178, 0.154]		-0.143 [-0.433, 0.020]	
$\Delta_k \log \text{ financial wealth}_t$		-0.199 [-0.460, 0.040]		-0.382 [-0.770, -0.149]
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
N	1,234	1,258	1,455	1,489

Transformations of the risky asset share—In Section (2.2), we linearized the relationship between the risky asset share and liquid wealth. Our linearization is not the only way in which one could linearize the relationship. For example, one could linearize in such a way that the dependent variable on the left-hand side of Equation (9) is the k -period difference of some transformation of the risky asset share. Here we consider the log and logit transformations and we check whether we obtain similar results if we use the k -period difference in the log risky asset share (Specification (1) in Table A.2) or logit risky asset share (Specification (2) in Table A.2) as the dependent variable in our regressions. As the table shows, we still obtain a slightly negative coefficient, as in Table 4, between two and three standard errors from zero. In addition, Specification (3) uses as the dependent variable a risky asset share that accounts for leverage. Instead of scaling stock holdings by the amount of liquid assets, we scale by liquid wealth (liquid assets minus non-mortgage debt, such as credit card debt, for example). As a result, the risky asset share can be bigger than one. Because there are some households

with liquid wealth close to zero or negative at time t , we have to discard observations with negative liquid wealth and we winsorize values of the risky asset share above 2. In the main part of the paper, we scale by liquid assets to avoid such truncation and winsorizing, but as Table A.2 shows the OLS results are very similar when we scale by liquid wealth instead of liquid assets. The coefficient on Δ_k log liquid wealth is significantly negative in the 1984-1999 sample period. Only for Specification (2) in the 1999-2003 subsample, the coefficient is higher than in our basic Specification in Table 4. For all three Specifications in Table A.2 we also re-ran our TSLS regressions with the risky asset share transformations and obtained similar results. Overall, the results are not systematically different if we choose a different linearization or take into account leverage.

Table A.2: Changes in the Proportion of Liquid Assets Invested in Risky Assets: Transformations of the risky asset share

	$k = 5$			$k = 2$		
	(1984-1999)			(1984-1999)		
	(1)	(2)	(3)	(1)	(2)	(3)
Δ_k log liquid wealth $_t$	-0.061 (0.023)	-0.077 (0.038)	-0.063 (0.010)	0.017 (0.036)	0.144 (0.069)	-0.089 (0.034)
Preference shifters	Y	Y	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y	Y	Y
Adj. R^2	0.00	0.05	0.05	0.01	0.00	0.02
N	1,234	1,234	1,234	1,455	1,455	1,455

One might suspect that the log transformation could have a stronger effect on the regressions that use changes in the proportion of financial wealth invested in risky asset, because the financial wealth risky asset share has some negative skewness, with observations being somewhat concentrated close to one (see Table 1). However, as Table A.3 shows, the transformations make little difference. The point estimates using the log transformation (Specification 1) and the logit transformation (Specification 2) are again negative, as in Table 5. Note that the logit transformation yields a lower number of observations here, because it excludes households with more than 100% invested in risky assets. Recall that the denominator of the risky asset share in this case is financial wealth (unlike the liquid assets risky asset share, which has liquid assets, not liquid wealth, as the denominator), which can be smaller than the amount of risky financial assets. In summary, these tests confirm that our results are robust to choosing a different linearization of our estimating equation.

Sampling weights—Our tests in the main part of the paper weight all observations equally (except for the summary statistics in Table 1), despite the fact that households with different characteristics have different sampling probabilities in the PSID. The

Table A.3: Changes in the Proportion of Financial Assets Invested in Risky Assets: Transformations of the risky asset share

	$k = 5$		$k = 2$	
	(1984-1999)		(1999-2003)	
	(1)	(2)	(1)	(2)
$\Delta_k \log \text{ financial wealth}_t$	-0.225	-0.362	-0.165	-0.140
	(0.036)	(0.097)	(0.037)	(0.099)
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
Adj. R^2	0.18	0.10	0.10	0.04
N	1,254	1,142	1,487	1,303

reason is that our model should apply to all households. Hence, for estimating habit effects it would be inefficient to use sample weights (see, e.g., the discussion in Deaton (1997), p. 70). If some household characteristics need to be controlled for, we do so by including them in the regression. But it turns out that we also obtain similar results if we weight households by the PSID sample weights, as shown in Table A.4.

Table A.4: Changes in the Proportion of Liquid Assets Invested in Risky Assets: Weighted with Sampling Weights

	$k = 5$			$k = 2$		
	(1984-1999)			(1999-2003)		
	OLS	OLS	TOLS	OLS	OLS	TOLS
$\Delta_k \log \text{ liquid wealth}_t$	-0.015	-0.010	0.024	0.027	0.021	-0.108
	(0.007)	(0.008)	(0.065)	(0.011)	(0.015)	(0.066)
Asset composition controls		Y			Y	
Preference shifters	Y	Y	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y	Y	Y
Adj. R^2	0.01	0.01	–	0.01	0.02	–
N	1,234	1,234	1,234	1,455	1,455	1,455

The results are almost identical to those in Table 4. The same is true for the regressions with the proportion of financial wealth invested in risky assets (not tabulated).

LAD regressions—To check whether our results might be driven by outliers, we run least-absolute deviation (LAD) regressions (median regressions). The results are shown in Table A.5, with bootstrap standard errors in parentheses. It is apparent that the estimates are virtually identical to our OLS estimates. Therefore, we conclude that our results are not driven by outliers.

Table A.5: Changes in the Proportion of Liquid Assets Invested in Risky Assets: Median regressions

	$k = 5$		$k = 2$	
	(1984-1999)		(1999-2003)	
$\Delta_k \log \text{liquid wealth}_t$	-0.010	-0.009	0.022	0.016
	(0.015)	(0.028)	(0.011)	(0.013)
Asset composition controls		Y		Y
Preference shifters	Y	Y	Y	Y
Life-cycle controls	Y	Y	Y	Y
Year-region FE	Y	Y	Y	Y
Pseudo R^2	0.03	0.03	0.02	0.02
N	1,234	1,234	1,455	1,455