

Andersen, Asger Lau; Johannesen, Niels; Sheridan, Adam

Working Paper

Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market

CEBI Working Paper Series, No. 11/21

Provided in Cooperation with:

Center for Economic Behavior and Inequality (CEBI), Department of Economics, University of Copenhagen

Suggested Citation: Andersen, Asger Lau; Johannesen, Niels; Sheridan, Adam (2021) : Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market, CEBI Working Paper Series, No. 11/21, University of Copenhagen, Department of Economics, Center for Economic Behavior and Inequality (CEBI), Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/258955>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CEBI WORKING PAPER SERIES

Working Paper 11/21

Dynamic Spending Responses to Wealth Shocks:
Evidence from Quasi-lotteries on the Stock Market

Asger Lau Andersen

Niels Johannesen

Adam Sheridan

ISSN 2596-44TX

CEBI

Department of Economics
University of Copenhagen
www.cebi.ku.dk

Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market*

Asger Lau ANDERSEN (University of Copenhagen and CEBI)

Niels JOHANNESSEN (University of Copenhagen and CEBI)

Adam SHERIDAN (University of Copenhagen and CEBI)

July 6, 2021

Abstract

How much and over what horizon do households adjust their consumption in response to stock market wealth shocks? We address these questions using granular data on spending and stock portfolios from a large bank and exploiting lottery-like variation in gains across investors with similar portfolio characteristics. Consistent with the permanent income hypothesis, spending responses to stock market gains are *immediate* and *persistent*. The responses cumulate to a marginal propensity to consume of around 4% over a one-year horizon. The estimates differ substantially by household liquidity, but not by financial attention, as measured by the frequency of account logins.

*Asger Lau Andersen: asger.lau.andersen@econ.ku.dk; Niels Johannesen: niels.johannesen@econ.ku.dk; Adam Sheridan: adam.sheridan@econ.ku.dk. We are extremely grateful to key employees at Danske Bank for their help. All individual data used in this analysis has been anonymized and no single customer can be traced in the data. All data processing has been conducted in accordance with the bank's strict data privacy guidelines. We gratefully acknowledge financial support from the Danish National Research Foundation, the Independent Research Fund Denmark and the Danish Innovation Fund. The activities of CEBI are financed by the Danish National Research Foundation, Grant DNRF134. Declarations of interest: none

1 Introduction

The stock market is volatile and thus an engine of both wealth creation and wealth destruction. The bust and boom in 2020 is a case in point. In the first quarter, when the global economy was shaken by the onset of the COVID-19 pandemic, the stock market portfolios of U.S. households shrank by \$6 trillion, which compares to their total net worth of around \$100 trillion. In the next three quarters, markets reversed and the same households recorded an increase in their stock market wealth of \$11 trillion.

The implications of such swings in financial wealth for the real economy depend crucially on *how much* they induce households to change their consumption and *over what horizon*. If consumption responses are large and swift, the stock market may amplify the business cycle in the real economy: a boom creates financial wealth by driving up stock prices, which in turn stimulates consumption and reinforces the boom. By contrast, if consumption responses are weak and sluggish, this feedback mechanism is much less relevant for the business cycle. This is important for policy makers deciding how strongly to lean against the stock market to stabilize aggregate demand (Cieslak and Vissing-Jørgensen, 2020) and for macro-economists seeking to integrate equity markets into quantitative models (Kaplan and Violante, 2018). Careful empirical analysis is warranted as competing theories of household behavior have diverging implications: The permanent income hypothesis predicts a small, but immediate and persistent increase in consumption in response to windfall gains (Friedman, 1957) whereas buffer-stock models imply a larger, but temporary increase (Carroll, 1997) and theories of behavioral inattention suggest a delayed response, if any at all (Gabaix, 2019).

In this paper, we combine rich data and a novel empirical strategy to estimate how shocks to stock market wealth affect consumer spending over the short and the medium term. Our main data source is customer records from Danske Bank, the largest retail bank in Denmark, which allow us to track the consumer spending and investment portfolios of almost 400.000 investors over eight years. On the spending side, we observe card purchases, bill payments, and cash withdrawals at the transaction level and aggregate them to construct monthly household-level measures of consumer spending.¹ On the investment side, we observe portfolios and stock prices each day and use them to measure gains and to summarize portfolio characteristics. Besides these main variables, the bank data also contains information on the frequency of

¹We follow a recent literature that uses transaction data from banks and financial apps to measure consumer spending (e.g. Gelman et al., 2013; Baker, 2018; Ganong and Noel, 2019; Sheridan et al., 2020). We note that spending is not always equivalent to consumption, e.g. durable goods enter our measure of spending in the month they are paid, but are typically consumed over a protracted period.

account logins, which allows us to explore the role of financial attention (Sicherman et al., 2016). Finally, we combine the bank data with administrative data from government registers on employment relations, bank relations, income, balance sheets, education, amongst other things. The administrative data indicates that the households in our sample generally hold few stocks outside Danske Bank, implying that our portfolio data are close to complete.

Our empirical strategy exploits the comprehensive data on stock market portfolios to isolate lottery-like variation in gains. Intuitively, we compare households who enter the month with stock market portfolios that are highly similar in terms of size, risk and expected returns, but earn different gains over the month because they consist of different stocks with idiosyncratic returns. Our key assumption is that the variation in gains across portfolios with highly similar characteristics is quasi-random and thus exogenous to the spending decisions of portfolio owners. In support of this assumption, we document that once we condition on portfolio characteristics, portfolio gains are orthogonal to a range of ex ante portfolio owner characteristics, including past levels of spending.²

Based on these considerations, we develop a model that identifies the dynamic spending response to stock market wealth shocks. The main explanatory variable is the household’s stock market gains over the current month and the dependent variable is the change in the household’s spending between the previous month and the current or some future month. By varying the time horizon over which changes in spending are measured, we estimate immediate as well as lagged responses.³ We modify this simple framework in two ways to ensure that the estimates are identified from plausibly random variation in gains. First, we instrument the household’s *actual* gains, which are endogenous to trading during the month, with the household’s *passive* gains, defined as the gains earned by the portfolio held at the beginning of the month. Second, we include granular controls for portfolio characteristics: a full set of interactions between indicators of portfolio value (size), volatility of past returns (risk), average of past returns (expected return), and month fixed effects. Thus, identification rests on the assumption that the instrument – the gains of the portfolio held at the beginning of the month – is randomly assigned to households conditional on the size, risk and expected return of this portfolio.

Our main estimates show that spending responses to stock market gains are *immediate* and *persistent*. Specifically, a \$1 gain is associated with an increase in spending of around 0.2 cents

²These tests are inspired by the recent literature using actual lotteries to estimate the causal effect of wealth on labor supply and child development (Cesarini et al., 2016, 2017).

³This is akin to the local projections approach widely used in macroeconomics (Jordà, 2005).

in the first month, around 4 cents over a one-year horizon, and around 12 cents over a three-year horizon. The estimated spending response is roughly uniform across months. The findings are consistent with the permanent income hypothesis where households consider gains in the stock market as windfalls and respond by adjusting spending by a small amount in the present and all future periods.

We probe the validity and robustness of the results in a number of ways. We estimate the model for negative horizons and find no effect of gains on *past* changes in spending. The parallel spending trajectories in the pre-period support our causal interpretation of the diverging spending trajectories in the post-period. We also re-estimate the model for various subsamples to address specific concerns about identification. First, using administrative data on bank relations, we restrict the sample to households who bank exclusively with Danske Bank to eliminate the potential bias created by spending and investments through other banks. Second, using administrative data on employment relations, we address concerns about confounding shocks to expected future wages. While we always exclude households investing in a firm that is also the primary employer of a household member, our robustness tests go further and disregard any investments in firms that are in the same industry as a primary employer.

We broaden the analysis to cover two outcomes that are important to understand household behavior and decision-making processes in the context of wealth shocks. We first show that stock market gains have an immediate effect on investment decisions: A \$1 gain is associated with a decrease in net investment of 2 cents in the first month and 10 cents over a one-year horizon. Disinvestment converts paper gains into cash dividends that can finance increased spending, but may also serve to rebalance the portfolio across safe and risky assets. We then show that stock market gains have a significant and lasting effect on financial attention as measured by account logins. While the immediacy of the consumption and investment responses suggest that households are generally attentive to the performance of their portfolios, this finding suggest that financial attention is itself endogenous to stock market gains.

Finally, we learn more about the mechanisms underlying the responses by investigating heterogeneity in key dimensions. The most striking pattern emerges when we split the sample by liquidity: A \$1 gain increases spending by around 15 cents over a one-year horizon for the bottom tercile, but only by 3 cents for the top tercile. By contrast, we find little heterogeneity when we split the sample by financial attention, measured as the frequency of logins. Even the least attentive households exhibit significant responses to stock market wealth shocks over a three-month horizon, both in terms of spending and investment, and the magnitude of their

responses is only marginally smaller than for the most attentive households. Thus, inattention does not appear to play an important role in dampening household responses to stock market cycles even over relatively short time horizons.

Our paper contributes in several ways to the literature that studies how stock markets affect household consumption through the wealth channel. Methodologically, we push the frontier by developing a design that takes advantage of highly granular data from financial institutions to deliver estimates of dynamic spending responses under transparent assumptions. Most micro-studies use much cruder data from household surveys and effectively identify the wealth effect by comparing the consumption of stockholders and non-stockholders (e.g. Mankiw and Zeldes, 1991) or stockholders with different portfolio sizes assuming that all households earn the market return (e.g. Dynan and Maki, 2001; Paiella and Pistaferri, 2017).⁴ As the identifying variation in gains in such studies derives from differences in the size and composition of wealth, it is likely endogenous to unobserved factors that also affect consumption decisions. Recent studies have made progress using better data and new empirical designs. Meyer, Pagel and Previtero (2019) and Bräuer, Hackethal and Hanspal (2020) use data similar to ours to investigate how spending responds to sharp changes in liquidity, due to forced realizations and dividend payouts respectively, but do not consider wealth shocks. Di Maggio, Kermani and Majlesi (2020) use administrative data on income and balance sheets to estimate the same-year effect of stock market gains on imputed consumption while focusing on the heterogeneous responses across the wealth distribution. They do not characterize the dynamics of the adjustment, whether it is immediate or sluggish, persistent or transient, which is a key focus in our paper.⁵

Our estimates speak directly to macro questions about the effect of asset prices on aggregate consumption where the key parameter is the average marginal propensity to consume across all stockholders, weighted by their stock market wealth. Since our dataset covers all customers in a nationally representative bank, our full-sample estimates naturally approximate this parameter. Our estimate that households spend around 4% of their aggregate stock market gains over a one-year horizon is somewhat higher than the estimate in a prominent recent study using aggregated data (Chodorow-Reich, Nenov, and Simsek, 2021) as well as the central scenario in an influential review paper (Poterba, 2000), both of which are around 3%. Given the magnitude

⁴Paiella and Pistaferri (2017) exploit the survey nature of their data to elicit expectations about returns, which uniquely allows them to distinguish between the consumption effects of anticipated and unanticipated wealth shocks.

⁵Another branch of the literature uses aggregate data on stock markets and consumption (e.g. Lettau and Ludvigson, 2004; Carroll, Otsuka and Slacalek, 2011). Most recently, Chodorow-Reich, Nenov and Simsek (2021) use county-level data to estimate the effect of stock market wealth on local employment outcomes and infer the underlying consumption dynamics from a quantitative model.

of stock market gains, even small differences in the estimated spending responses have significant implications for the strength of the spillovers to the real economy.⁶

Our results also inform theories about household consumption behavior. A central theme is that consumption, contrary to the predictions of standard theory, may adjust to shocks in a sluggish fashion because of habits (Campbell and Cochrane, 1999), inattention (Reis, 2006) or fixed consumption commitments (Chetty and Szeidl, 2016). Strikingly, our estimates of the dynamic spending response to stock market gains exhibit little excess smoothness: stock market wealth shocks are accompanied by a persistent adjustment of the level of spending in the same month. We consistently find this pattern across all the subsamples we are studying and we find no meaningful differences across investors with different levels of financial attention. The results suggest that stockholders generally pay sufficient attention to stock prices to account for wealth shocks in their financial planning.

Relatedly, we contribute to the literature on heterogeneity in the marginal propensity to consume, which is central to a new generation of macro models with heterogeneous agents (Kaplan and Violante, 2014; Kaplan, Moll and Violante, 2018; Auclert, 2019; Auclert, Rognlie and Straub, 2020). The striking heterogeneity by liquidity is notable given that our sample consists of stockowners who are relatively liquid. It resonates with existing work on stock market wealth shocks (Di Maggio, Kermani, and Majlesi, 2020) as well as studies of spending responses to labor income shocks (Ganong et al., 2020), stimulus payments (Johnson, Parker and Souleles, 2006), unemployment shocks (Andersen et al., 2020) and lottery gains (Fagereng, Holm and Natvik, 2021).⁷

Finally, we make several contributions to the broader literature on household financial behavior. First, our finding that households make active investment decisions that mitigate the changes in portfolio composition created by capital gains and losses represents new evidence on portfolio rebalancing (Calvet, Campbell and Sodini, 2009). Our investment results are identified from the same lottery-like variation as the spending results, which ensures that they are not confounded by, for instance, life-cycle patterns in portfolio choices (Fagereng, Gottlieb and Guiso, 2017). Second, we contribute new evidence on selective financial attention. A theoretic-

⁶The range of estimates reported in the literature is wide. Some find that households consume 10-15% of stock market gains within a year (Dynan and Maki, 2001) while others report estimates close to zero (Paiella and Pistaferri, 2017). Di Maggio, Kermani and Majlesi (2020) do not report a comparable estimate for the full sample, but their estimates for subsamples range from more than 20% at the bottom of the wealth distribution to below 3% for the top 10%. In the broader literature on wealth effects, recent estimates of the marginal propensity to consume out of housing wealth shocks are 5-7% (Mian, Rao and Sufi, 2013) and 5% (Aladangady, 2017).

⁷A notable exception is Kueng (2018) who finds that spending responses to payments from the Alaska Permanent Fund are driven by high-liquidity households.

cal literature posits that information and beliefs affect utility directly (e.g. Brunnermeier and Parker, 2005), implying that individuals may choose not to update their information set even when it is costless. An earlier application to the stock market documents that investors pay less attention to their financial accounts around trading days where markets are down and the risk of receiving bad news is high (Sicherman et al., 2016). We improve the identification of this behavioral mechanism, as our empirical framework absorbs the effect of confounding events that affect both market returns and average financial attention.⁸ We also document that the effect of stock market wealth shocks on financial attention is highly persistent and remains detectable after as long as two years.

The paper proceeds in the following way. Section 2 describes the data. Section 3 develops the empirical framework. Sections 4-6 present the results. Section 5 concludes.

2 Data

Our primary data source is the complete customer records of all customers at Danske Bank, the largest retail bank in Denmark, for the period 2009-2016. We draw on this data to construct the main variables: consumer spending, stock market gains, portfolio characteristics and the frequency of online account logins. We add administrative data from various government registers, which serves a number of auxiliary purposes such as forming households, identifying employers, and assessing the completeness of the bank data. All data sources identify individuals by their unique personal identification number and can thus be seamlessly merged.

2.1 Bank data

The bank data includes transaction-level records for all bank accounts held by customers at Danske Bank, which we use to construct a measure of *consumer spending* at the monthly frequency.⁹ The spending measure aggregates purchases made with debit and credit cards, bill payments, in-store mobile payments and cash withdrawals. Using the international standard for classification of merchants (Merchant Category Codes) and a proprietary register categorizing bill creditors, we exclude transactions that are not associated with consumption, such as debt service and tax payments. Moreover, to ensure comparability across home owners and renters, we also exclude transactions related to housing expenditure, such as rent payments. In other

⁸For instance, wars, disasters and pandemics may at the same time depress market returns and crowd out attention to financial accounts.

⁹The use of transaction-level data from banks and financial apps to measure consumer spending was pioneered by Gelman et al. (2013) and validated against other data sources by Baker (2018).

work drawing on the same transaction data, we show that both the level and the dynamics of our spending measures are consistent with other data sources (Andersen et al., 2020).¹⁰

The bank data also includes daily information about assets and prices, the same information the bank uses to compute actual, customer-facing measures of portfolio values and returns. We delimit a household’s stock market portfolio to include direct equity investments and mutual fund shares, but not direct bond investments and deposits.¹¹ Moreover, we include all liquid securities accounts, but exclude retirement savings accounts where assets typically cannot be liquidated without triggering significant tax penalties.

We summarize the portfolio information in the variable $a_{i,j,d}$, stating the number of security j held by household i on day d , and the price information in the variable $\Delta p_{j,d}$, stating the change in the price of security j on day d adjusted for stock splits, mergers of assets, dividend payouts and other corporate events. With this notation, we define stock market gains in month t in the following way:

$$G_{i,t} = \sum_{d \in t} \sum_j \Delta p_{j,d} \cdot a_{i,j,d} \quad (1)$$

where gains are summed over days of the month and over the individual securities in the household’s stock market portfolio. As $\Delta p_{j,d}$ is the price change adjusted for payouts, $G_{i,t}$ includes capital gains as well as dividends. Our empirical strategy isolates quasi-random variation in $G_{i,t}$, but not in capital gains and dividends separately; hence, we do not distinguish between the two forms of gains in the analysis.¹² We note that the intra-day variation in prices may introduce a small error in the measurement of gains on days where a particular asset is traded.

A key element in the empirical strategy is the controls for ex ante portfolio characteristics. Using the bank data on portfolios and prices, we measure *portfolio size* in month t as the combined value of stocks in the portfolio on the last day of the previous month; *risk* as the standard deviation of this portfolio’s returns over the past twelve months; and *expected returns* as the mean of this portfolio’s returns over the past twelve months where the portfolio return is itself defined as gains in a given month measured relative to the portfolio size.

¹⁰Specifically, we show that overall spending levels in our data line up well with the levels observed in the Consumer Expenditure Survey from Statistics Denmark and that the monthly dynamics in card spending for the full sample of Danske Bank customers mirrors publicly available statistics on aggregate card spending almost perfectly (Andersen et al., 2020).

¹¹This is consistent with the risky portfolio studied by Calvet, Campbell and Sodini (2007, 2009).

¹²Several recent papers have estimated responses to dividend payouts (Bräuer, Hackethal and Hanspal, 2020) and, relatedly, forced realizations (Meyer, Pagel and Previtero, 2019) employing empirical approaches specifically designed to isolate quasi-random variation in these flows. Other papers compare household responses to capital gains and dividends (Di Maggio, Kermani and Majlesi, 2020).

We also use the portfolio data to construct a monthly measure of net investment in the stock market, that is gross purchases of stocks net of gross sales. It is convenient to compute net investment residually as the total change in the value of the household’s stock market portfolio net of the change that is due to capital gains.

Finally, we create a measure of financial attention by counting the number of days in a month that households log into their Danske Bank accounts. Following Karlsson, Loewenstein and Seppi (2009), a number of recent studies have used account logins to capture households’ attention to private finances empirically (Sicherman et al., 2016; Olafsson and Pagel, 2017; Gargano and Rossi, 2018). While most of these studies obtain data from investment brokers and therefore only include logins to brokerage accounts, our measure is broader by counting logins to all types of customer accounts.

2.2 Government register data

We merge the customer data from Danske Bank with administrative data from various government registers. Specifically, we add information about demographics and household identifiers from the *population register*; about income and financial accounts from the *tax register*; about employment relations from the *employment register*; and about schooling from the *education register*. This enables us to address a range of concerns about identification.

The population register includes household identifiers that link cohabiting couples and dependants. This allows us to aggregate both spending, income and stock market variables across members of the same household and estimate the marginal propensity to consume at the household level. This is important because not all couples have shared ownership of all financial accounts. To the extent that the household’s securities are nominally owned by one spouse, but gains feed into the spending of both spouses, conducting the analysis at the individual level would understate the spending response.

The tax register includes compulsory account-level reports filed by all Danish banks to the tax authorities at the end of each year. The reports help us address concerns about the completeness of the bank data, which are inherent to studies using data from a single bank or financial app (Baker, 2018). First, the reports allow us to measure the combined value of stock market portfolios held in all banks for our sample of Danske Bank customers and compare it to the value of stock market portfolios that we observe in the data from Danske Bank. As shown in Figure 1A, the correspondence between portfolios in Danske Bank and portfolios overall is generally close to perfect, implying that the households in our sample hold very few

stocks in other banks. There is some divergence at the top suggesting that large investors more frequently use multiple banks, but even for the very largest investors less than 20% of the combined portfolios is unobserved in the granular data from Danske Bank. Second, we use the bank reports to the tax authorities to construct a subsample of households who are *exclusive* customers at Danske Bank in the sense that they have no accounts at other banks at the beginning nor at the end of the year. In robustness checks, we re-estimate the model for this subsample to test whether our estimates are sensitive to incomplete coverage.

The tax register also includes information on income, which is both reliable and comprehensive, as it is largely derived from reports by third parties such as employers and financial institutions rather than by taxpayers themselves (Kleven et al., 2011; Alstadsæter et al., 2019).¹³ Comparing to total after-tax income allows for a useful check of our transaction-based measure of spending. As shown in Figure 1B, spending increases monotonically with income in the cross-section with a gradient less than unity that reflects the transitory component of income. Many households in the sample spend more than their income, which is not surprising given the demographic composition with a high fraction of elderly (see below).

The employment register contains links between individuals and their employers. We create a crosswalk between employer identifiers in the employment register and security identifiers in the portfolio data (ISIN) and use it to identify households who invest in a firm or, more broadly, in an industry where they are also employed. This helps us address the concern that the spending of employees-shareholders may be affected by shocks to firm profitability not just through stock market gains, but also through changes in expected future wages. As there is no obvious way to disentangle the two effects, our analysis generally excludes employees-shareholders from the sample: We drop all households who hold investments in a firm that is also the primary employer of a household member. Further, we conduct robustness tests where we address the possibility that wage expectations correlate within industries by disregarding all investments in firms that belong to the same industry as a primary employer.

Finally, we construct variables that serve as household-level controls based on information from the administrative registers. In the education register, we observe the schooling and degrees obtained by each individual and use it to construct categorical variables indicating the highest level of education in the household: primary school, high school, college and graduate degree. In the population register, we observe the age of each individual and use it to construct a variable

¹³An alternative measure of income can be constructed from the transaction-level bank data by employing an algorithm that classifies incoming transfers as either income or other transfers (Sheridan et al., 2020); however, such as measure will generally be noisier than the income information from administrative tax data.

indicating the age of the oldest household member. We also construct a measure of household size, indicating the number of household members.

2.3 Sample and descriptive statistics

Starting from the sample of all households with securities accounts at Danske Bank, we impose a number of mild restrictions to arrive at our main estimating sample. First, we require that households have non-negligible stock market portfolios and thus exclude households with portfolios of less than 100 kroner (around \$15). Second, we require that households are active Danske Bank customers in the sense that they have at least five spending transactions in a given month. Third, we drop households who invest in a firm that is also the primary employer of one household member because of the potentially confounding effect of expected future wages (discussed above). Finally, we drop households whose only stock market investment is shares in Danske Bank. With these sample restrictions, we arrive at our estimation sample with around 13.5 million household-month observations comprising 390,000 distinct households.

Table 1 reports descriptive statistics for the estimation sample. The most salient demographic characteristic is *age*: the oldest member of the average household is more than 60 years old. This is consistent with the empirical regularity that stock market participation tends to be low early in the life-cycle (e.g. Gomes and Michaelides, 2005). At the sample mean, the value of the stock market portfolio is around \$80,000, which is close to half of the household's financial wealth. The average monthly stock market gain is around \$600, which is substantial compared to monthly income of \$6,500 and monthly spending of around \$3,600. Gains largely reflect price changes and only to a limited extent dividends.

Figure 2 describes the cross-sectional differences in portfolio returns, which is ultimately the variation that identifies our empirical results. One may be concerned that stockholders in each month may have very similar returns, either because they hold similar portfolios, e.g. through broad mutual funds, or because the returns of individual stocks are highly correlated. Figure 2 shows that the mean portfolio return across the stockholders in our sample (orange line) is generally less volatile than the main Danish stock market index (black line), which is consistent with diversification through foreign stock investments. However, in each month, there is significant variation in portfolio returns across stockholders, as illustrated by the wide range between the 5th and 95 percentile of the portfolio returns (shaded area). This range is generally larger than 10 percentage points and in periods with market turmoil, e.g. in early 2009 following the financial crisis, as large as 40 percentage points.

3 Empirical framework

Our objective is to estimate how much and over what time horizon households change spending in response to stock market gains. In this section, we develop a simple empirical framework that identifies the dynamic spending responses from lottery-like variation in gains deriving from variation in returns across portfolios with similar characteristics. We also report results from explicit tests of the quasi-random nature of the identifying variation adopted from the literature on lotteries (e.g. Cesarini et al., 2016, 2017).

3.1 Baseline model

We denote consumption by $C_{i,t}$ and wealth by $W_{i,t}$ where i refers to households and t refers to the month. We are interested in the causal effect of stock market gains $G_{i,t}$ on consumption. As gains represent a change in wealth, this is conceptually the wealth effect on consumption (Paiella and Pistaferri, 2017). Drawing on this notation, we can formulate the following naive model:¹⁴

$$C_{i,t+1} - C_{i,t} = \alpha + \beta G_{i,t+1} + \epsilon_{i,t+1} \quad (2)$$

where β expresses the marginal propensity to consume out of wealth gains. The naive model suffers from potentially severe endogeneity problems because stock market gains are not randomly assigned. Although gains in a given month have a random component due to the unpredictable nature of returns, they are also endogenous with respect to two dimensions of household choice: *ex ante* decisions about portfolio characteristics made before the month begins and *ex post* decisions about sales and purchases made in the course of the month.

First, a household's gains in a given period depend on the characteristics of its *ex ante* portfolio. Everything else equal, larger portfolios are associated with larger gains because the portfolio return applies to more stocks; riskier portfolios are associated with larger gains and larger losses because the portfolio return is more volatile; and portfolios with higher expected returns are associated with larger gains because the portfolio return has a higher mean. To the extent that households holding larger, riskier and higher-yield portfolios are systematically different with respect to their spending dynamics, the estimate of β may be biased.¹⁵

¹⁴We follow most papers in the literature in relating ΔC to ΔW (e.g. Dynan and Maki, 2001; Paiella and Pistaferri, 2017; Aladangady, 2017). Some papers, however, relate $\Delta \log(C)$ to $\Delta \log(W)$ and thus estimate elasticities rather than MPCs (e.g. Mian, Rao and Sufi, 2013) while others relate ΔC to ΔG (e.g. Di Maggio et al., 2020).

¹⁵According to standard portfolio theory, the expected return is positively related to risk; however, some

We address this type of endogeneity by controlling for the portfolio characteristics with a rich set of non-parametric controls. Concretely, we create 100 indicators of *portfolio size* based on a ranking of households by the value of their portfolio at the beginning of the month; 20 indicators of *portfolio risk* based on a ranking of households by the standard deviation of the return of this portfolio over the past twelve months; and 20 indicators of *expected return* based on a ranking of households by the mean return of this portfolio over the past twelve months. We augment the model with the term $\Theta_{i,t}$, which includes all combinations of these indicators ($100 \times 20 \times 20$ variables) interacted with month fixed effects. Hence, our model effectively compares the spending behavior of households who have different stock market gains in a given month despite entering the month with portfolios that have similar characteristics. The variation in gains arises because portfolios with similar characteristics are generally composed of different individual stocks with different idiosyncratic returns.

Second, a household's gains in a given period are affected by *ex post* decisions about buying and selling. For instance, two households who enter a period with exactly the same portfolio generally have different gains if one of them decides to liquidate the portfolio during the period and the other one decides to hold it throughout. This introduces an endogeneity problem as buying and selling may be shaped by a range of shocks that also influence spending decisions. It may also take the form of reverse causality if households sell stocks to raise cash for consumption.

Following other papers in the literature (e.g. Di Maggio et al., 2020), we address this problem by instrumenting, $G_{i,t+1}$, a household's *actual* gains in month $t + 1$, with $\bar{G}_{i,t+1}$, the household's *passive* gains in month $t + 1$, defined as the gains of the portfolio that the household held at the beginning of the month. Passive gains thus express the counterfactual stock market gain that the household would have received, absent portfolio adjustments. For households that do not trade during month $t+1$, the counterfactual gain coincides with the actual gain. Since households generally trade only a small fraction of their portfolio in a given month, the instrument is highly relevant.¹⁶

3.2 Identifying assumption

Once we control for portfolio characteristics and instrument actual gains, the key identifying assumption of the model is that the instrument, gains of the portfolio held at the beginning

households are underdiversified and therefore not at the efficient frontier (Calvet, Campbell and Sodini 2007). This implies that there is variation in expected returns across portfolios with the same risk that are at different distance to the frontier.

¹⁶Concretely, we obtain a highly significant coefficient of 0.986 when we regress actual gains on passive gains (and the full set of controls) in the estimation sample.

of the month, is randomly assigned to households conditional on the size of this portfolio and the mean and the standard deviation of its past returns. This assumption is consistent with a simple theoretical model of portfolio choice where investors have preferences over the mean and the variance of future returns (Markowitz, 1952) and form beliefs about these moments from past returns. Such investors are *ex ante* indifferent between portfolios whose past returns exhibit the same combination of mean and variance and *ex post* differences in realized returns are therefore orthogonal to investor characteristics conditional on these portfolio characteristics. By contrast, if investors have preferences over other moments of the return distribution (Harvey and Siddique, 2000) or differ systematically in their ability to obtain or process information (Fagereng et al., 2020), it could pose a risk to our identification strategy. Specifically, it is conceivable that investors with privileged access to information or superior cognitive skills are able to pick portfolios that perform better conditional on past returns, in which case the assignment of the instrument would not be conditionally random.

To assess whether this threat to identification is significant, we conduct a range of tests adopted from the literature on lotteries (e.g. Cesarini et al., 2016, 2017). We note that the identifying assumption has the stark implication that passive gains should be uncorrelated with *any* *ex ante* investor characteristic conditional on portfolio controls. We test this implication directly for five investor characteristics that are observable in our dataset and report the results in Figure 3. Each panel displays a binned scatterplot of passive gains against an investor characteristic, unconditionally (blue dots) as well as conditional on the portfolio controls (orange diamonds). The results document a strong unconditional correlation between investor characteristics and passive gains: investors who are older, hold more deposits, have higher monthly income and spending and invest a larger share of their financial wealth in the stock market tend to have larger stock market gains in the average month. However, consistent with our identifying assumption, these correlations all vanish almost completely once we introduce portfolio controls.

3.3 Dynamics

We generalize the model to capture the full dynamic spending response to stock market gains. To achieve that, we consider the change in spending between month $t - 1$ and $t + h$ where h indexes the horizon, positive or negative, over which spending responses are measured.¹⁷ We thus estimate the following model separately for each different value of h :

¹⁷This is similar to the local projections methodology widely used in macroeconomics (Jorda, 2005).

$$C_{i,t+h} - C_{i,t} = \alpha + \beta G_{i,t+1} + \Theta_{i,t} + \gamma X_{i,t} + \mu_{i,t+h} \quad (3)$$

where the actual stock market gains are instrumented with the passive gains. The vector $\Theta_{i,t}$ is the set of ex ante portfolio characteristics described above (size \times risk \times expected return). The vector $X_{i,t}$ is a set of socio-demographic household characteristics included to reduce the residual variation: the age of the household's oldest member, the number of children in the household, and the highest level of education completed by a household member.

For $h = 1$, the dependent variable is simply the one-period change in spending and the model thus estimates the effect of a stock market gain in a given month on spending in the same month. This is conceptually similar to most of the estimates in the literature, yet different because our data has a higher frequency: with monthly observations, our estimate has the flavor of an instantaneous spending response.

For $h > 1$, we trace out the lagged spending response to stock market gains: the effect of a stock market gain in a given month on spending one, two, three or more months later. Lagged responses arise in many theoretical models, including the canonical permanent income hypothesis according to which households use windfall gains to increase consumption in all future periods (Friedman, 1957).

For $h < 0$, we estimate how gains correlate with *past* spending changes, an important diagnostic of endogeneity. Our key identifying assumption that passive gains are conditionally random implies parallel trends in the pre-period: investors who enter a month with similar portfolios should not have systematically different spending trajectories in the past depending on their passive gains in that month. Diverging spending trajectories in the pre-period would be indicative that the model suffers from endogeneity problems.

When we generalize the model to capture the dynamic responses to gains, the autocorrelation of the passive gains becomes crucial for the interpretation. If there is no autocorrelation, investors with different passive gains in month $t + 1$ have the same expected passive gains in all other months. In that case, our estimates of how gains in month $t + 1$ affect spending in month $t + h$ are not confounded by differential gains in month $t + 2$, $t + 3$ and so on. However, if the autocorrelation is not zero, our dynamic estimates may pick up not just the effects of gains in month $t + 1$, but also the effect of differential gains in other months. While the efficient market hypothesis establishes a strong theoretical prior that stock returns are uncorrelated over time, it has been challenged by empirical work on return momentum (Fama and French, 2012).

To assess the empirical relevance of this concern, we estimate equation (3) using the change

in the value of the portfolio held at the beginning of month $t + 1$ as the dependent variable.¹⁸ Figure 4 illustrates the resulting estimates (blue line). Mechanically, the estimate for month $t + 1$ is unity: a \$1 gain immediately raises the value of the fixed portfolio by the same amount. In other months, the estimates are hovering around zero for negative horizons while they are above unity and increasing slightly over time for positive horizons. The results suggest that there is no autocorrelation in passive gains beyond the effect of compounding: gains in $t + 1$ do not predict gains in earlier months and only predict gains in later months because they themselves earn gains. To illustrate the latter point, we add another line to the figure: the effect of a \$1 gain in month $t + 1$ on the value of a fixed portfolio assuming that the portfolio, including the gain itself, earns the sample average of portfolio returns in every month (orange dashed line). Our estimates generally line up well with this stylized pattern, suggesting that compounding is the only source of autocorrelation in passive gains.

4 Results: Spending

This section presents our findings on how consumer spending responds to stock market gains. We first illustrate the dynamic spending responses and report the marginal propensity to consume cumulated over different horizons. Then, we show the results from a range of robustness tests. Finally, we document how spending responses vary in four dimensions of heterogeneity: liquidity, age, returns and financial attention.

4.1 Dynamic spending responses

We illustrate the estimated dynamic spending responses to a \$1 dollar gain in Figure 5. The gain causes a significant increase in spending of around 0.2 cents already in the same month and similar, or slightly larger, increases in spending in all subsequent months. Consistent with the identifying assumption that gains are randomly assigned conditional on the controls, there are no systematic differences in spending trajectories in the pre-period. Our empirical model thus compares investors who are on parallel spending trajectories before being exposed to a differential stock market wealth shock. Those who gain more increase spending differentially precisely in the month of the shock and the difference remains significant and roughly constant throughout the estimation window.

¹⁸More precisely, for each household and each month, we construct a hypothetical series of passive past and future portfolio values by keeping the composition of the current portfolio fixed and applying the actual past and future price changes. The passive values coincide with actual future portfolio values except when households trade. We then use the change in the passive value from month t to $t + h$ as the dependent variable.

We cumulate the marginal propensities to consume out of stock market gains over different time horizons and report the results in Table 2.¹⁹ After 6 months, investors have spent an additional 1.6 cents for each \$1 gain, increasing to 3.8 cents after one year, 7.5 cents after two years, and 11.9 cents after three years. Hence, today’s stock market gains appear to feed into future household consumption at a relatively constant rate of approximately 4% per year.

These findings are consistent with the permanent income hypothesis where households increase consumption by a small amount in the present and all future periods in response to windfall gains. Given the estimated marginal propensity to consume of around 4% and considering that stock market gains compound, the level shift in spending at the time of the gain may in fact be permanent and include future generations. This contrasts with the buffer-stock behavior found in the context of lottery prizes (Fagereng, Holm and Natvik, 2020) and unexpected inheritances (Druehl and Martinello, 2020) where households exhibit a much larger marginal propensity to consume and revert to the initial level of wealth. A plausible explanation for this discrepancy is differences in household characteristics. In particular, the average owner of stock market wealth, unlike the average lottery winner and the average recipient of an unexpected inheritance, holds high levels of liquid assets is not liquidity constrained in the way that generates buffer-stock behavior. We further explore the role of liquidity in the heterogeneity analysis below.

4.2 Robustness

We probe the robustness of the results by modifying the estimating equation and the sample in a number of ways and illustrate the resulting dynamic estimates in Figure 6. First, we drop the requirement that households hold no investments in the firms where they are employed (“Full”). Second, we restrict the sample to households who are exclusive customers at Danske Bank to address the concern that spending through accounts in other banks may confound the estimates (“Exclusive”).²⁰ Third, we ignore investments in firms belonging to the same industry as the employer of a household member when defining stock market portfolios in order

¹⁹We estimate cumulative spending effects over T months by using $(C_{i,t+1}-C_{i,t})+(C_{i,t+2}-C_{i,t})+\dots+(C_{i,t+T}-C_{i,t})$ as dependent variable in the baseline model. The estimates are not identical to the sum of the monthly coefficients from the dynamic model as it involves a more demanding sample requirement. Specifically, when estimating cumulative spending effects over T months, the dependent variable is only defined when households meet the sample requirement in every month between t and T . By contrast, when estimating the effect in month τ , the dependent variable is defined as long as households meet the sample requirement in month t and month τ .

²⁰Around half of the households in the baseline sample are exclusive customers at Danske Bank, as reported in Table 1.

to eliminate confounding effects through changes in wage expectations ("Excluding own industry").²¹ Fourth, we estimate the model without the household-level controls age, household size and education ("No controls"). Finally, we exclude observations where investors hold unlisted shares to address the concern that unobserved gains on unlisted shares may correlate, positively or negatively, with observed gains on listed shares and thus introduce an omitted variable bias ("Excluding unlisted").²²

In all of these alternative samples and specifications, the estimated dynamics of the spending response are similar to the baseline with roughly parallel trends in the pre-period and a pronounced differential level shift in the month of the differential wealth shock. The estimates of the cumulative spending responses are also comparable to the baseline, as shown in Table 2. The largest estimates emerge when we restrict the sample to exclusive Danske Bank customers (5.4 cents vs 3.8 cents in the baseline over a one-year horizon). This may potentially reflect that the baseline estimates fail to capture spending responses through unobserved accounts in other banks; however, the estimates for the smaller sample of exclusive Danske Bank customers are noisier and statistically indistinguishable from the baseline.²³

4.3 Heterogeneity

We investigate the heterogeneity in spending responses in four dimensions: liquidity, age, return and attention. For each household-month observation, we measure *liquidity* as liquid assets (deposits, stocks and other securities) held at the beginning of the month relative to average monthly income; *household age* as the age of the oldest household member at the beginning of the month; *absolute return* as the absolute size of the stock market portfolio return in the month; and *financial attention* as the number of days that a member of the household logged into their Danske Bank accounts in an average month.

In each dimension, we rank the observations within each month and split them into three equally sized groups. Table 3 summarizes the heterogeneity in each dimension. The mean liquidity ratio ranges from around 3 in the bottom group to more than 70 in the top group. In the bottom group, liquid assets are below two months of income for around 40% of the observations, a frequently used indicator of low liquidity (Zeldes, 1989). The attention variable

²¹Around 8% of the households in our baseline sample have investments in firms belonging to the same industry as their employer and such investments account for around 4% of the aggregate stock market portfolio, as reported in Table 1.

²²Around 10% of the households in our baseline sample own unlisted shares, as reported in Table 1.

²³Another possible explanation is selection into the sample of exclusive Danske Bank customers. Households with multiple bank relations typically have higher incomes and larger stock market portfolios, which may at the same time be associated with a lower marginal propensity to consume (see heterogeneity analysis below).

ranges from around 1 in the least attentive group, implying that these households log into their accounts on average one day per month, to around 12 in the most attentive group. The average household age is around 40 years in the youngest group compared to almost 80 years in the oldest group. The mean of the absolute portfolio return ranges from 0.6% in the lowest group to 7.7% in the highest group.

Figure 7 shows short-run responses (cumulated over 3 months) and medium-run responses (cumulated over one and two years) for each group. Technically, we obtain the estimates from a modified version of the baseline model where gains are interacted with three indicators capturing the heterogeneity in the dimension of interest.²⁴

When we split the sample by *liquidity*, the responses exhibit a clear gradient: less liquid households spend a larger share of their gains both in the short and the medium run. Specifically, the subsample with the lowest liquidity spend almost 15 cents of a \$1 gain over a one-year horizon and around 30 cents over a two-year horizon, around five times more than the subsample with the highest liquidity. This result resonates with a range of models highlighting the role of liquidity in shaping consumption responses to income and wealth shocks (e.g. Carroll 1997; Kaplan and Violante, 2014) and with recent empirical studies of the wealth effect on consumption (Di Maggio et al., 2020). While the result implies that a large group of stockowners with low liquidity consume their gains relatively quickly, this group hold small portfolios and therefore account for a modest share of the aggregate stock market wealth.

We find much less heterogeneity when we split the sample by *financial attention*. Even the least attentive households exhibit statistically significant responses to stock market wealth shocks over a three-month horizon and the magnitude of their response is only marginally smaller than for the most attentive households. Also over longer horizons, the heterogeneity in spending responses by financial attention is immaterial. The results suggest that financial inattention plays a limited role in attenuating household responses to stock market cycles even over relatively short time horizons.

Similarly, we find limited heterogeneity when we split the sample by *household age*. Contrary to the prediction of standard life-cycle models where older households consume more of their gains because their remaining life span is shorter and uncertainty about life-time income has largely resolved, we find that younger households consume slightly more, both in the short and the medium run. Several empirical papers have found a qualitatively similar relation between age and the marginal propensity to consume outside the domain of stock markets (e.g. Fagereng,

²⁴We show the full dynamics of the spending responses by liquidity, age, return and attention in Figure A1 in the Appendix.

Holm and Natvik, 2021).

Finally, allowing the estimates to vary with the *absolute return* reveals that the marginal propensity to consume stock market gains is significantly larger for returns close to zero than for returns in the tails of the distribution. A possible interpretation of this result is that the welfare cost of failing to smooth a wealth gain or loss over time is increasing in the absolute size of the gain or loss and thus, for a given portfolio size, increasing in the absolute return. This interpretation is reminiscent of recent work on excess sensitivity showing that predetermined income transfers trigger larger spending responses in the hands of high-income households for whom this departure from the optimal path is associated with smaller welfare losses (Kuong, 2018).

5 Results: Investment

In this section, we present our findings on how stock market gains affect net investment in the stock market. Net investment is directly linked to spending through the budget constraint: selling part of the portfolio is a way to convert paper gains into cash dividends that can finance spending.²⁵

We illustrate how net investment adjusts dynamically to stock market gains in Figure 8a. A gain of \$1 is associated with a significant disinvestment of more than 2 cents already in the same month and more modest responses in the following months. In the pre-period, net investment follows parallel trajectories across households with different gains, which is consistent with a causal interpretation of the divergence at the time of the differential wealth shock. The cumulative disinvestment response to a \$1 gain amounts to 14 cents over a one-year horizon, as shown in Table 4.

The results are consistent with investors creating their own dividends to finance spending out of capital gains. However, comparing the magnitude of the disinvestment and spending responses (e.g. -14 cents vs +4 cents over a one-year horizon), suggests that disinvestment also serves to rebalance the portfolio across safe and risky assets. This interpretation is consistent with earlier evidence that households make active investment decisions to, at least partially, offset the changes in portfolio composition created by capital gains (Calvet, Campbell and Sodini, 2009).

²⁵Such liquidations are of course not necessary for households to spend out of their stock market gains. They may choose to hold their portfolios fixed and finance spending by drawing down other liquid assets (e.g. deposit balances) or by borrowing (e.g. through credit cards).

We conduct the same battery of robustness tests as for spending and report the results in Table A1 in the Appendix. The baseline estimates are highly robust over shorter horizons, but become more sensitive to the specific choices as the horizon grows longer.²⁶ However, the qualitative conclusion that households disinvest significantly in response to stock market gains and that the disinvestment response exceeds the spending response by a wide margin stands in all the specifications.

Turning to heterogeneity, we show that investment responses exhibit a pronounced gradient in *liquidity*, notably in the short run, as shown in Figure 9.²⁷ The group with the lowest liquidity disinvest more than 12 cents of a \$1 gain over the first three months, around four times more than the group with the highest liquidity. This pattern mirrors the liquidity gradient in spending responses, which reinforces the notion that the two behavioral margins are intrinsically connected: Households with low liquidity have a stronger preference for converting capital gains into additional spending and, at the same time, have a greater need to sell stocks to finance an increase in spending in the short run. In the medium run, the differences in investment responses by liquidity are less stark.

The investment responses are much more weakly associated with *financial attention*. Although more attentive households disinvest slightly more in response to stock market gains, even the least attentive households exhibit statistically significant investment responses in the short run. The attention gradient appears to become somewhat steeper over time and for the least attentive group the effect on net investment is insignificant and close to zero over a two-year period.

Finally, *age* is a strong correlate of investment responses, with younger households disinvesting markedly more following gains, whereas the *absolute return* does not exhibit a clear monotonic association with investment responses.

6 Results: Attention

This section reports our results on the effect of stock market wealth shocks on financial attention. Behavioral theory highlights that attention may be selective and, more specifically, that investors may devote more attention to their finances after positive news than after negative news as

²⁶At one extreme, the cumulative disinvestment response to a \$1 gain over a one-year horizon is around 21 cents when we restrict the sample to exclusive Danske Bank customers while, at the other extreme, it drops to 13 cents when we include households with investments in firms that are also the main employer of a household member.

²⁷We show the full dynamics of the investment responses by liquidity, age, return and attention in Figure A2 in the Appendix.

attention amplifies the hedonic effect of new information (Karlsson, Loewenstein and Seppi, 2009). While our own heterogeneity analysis, presented in the previous sections, suggests that attention does not play an important role in mediating spending responses to wealth shocks, other work has shown that the attention of investors matters for individual outcomes, such as portfolio returns (Gargano and Rossi, 2018), as well as market outcomes, such as stock return volatility and risk premia (Andrei and Hasler, 2015).

We illustrate how the monthly number of account logins responds to stock market gains in Figure 8b. The variables are scaled so the coefficients indicate the responses to a gain of \$1,000. The results indicate that differential gains are associated with a sharp differential increase in attention that remains detectable throughout the estimation window. Specifically, a gain of \$1,000 causes 0.015 more logins in the first month and between 0.005 and 0.015 more logins in the following months. Over a one-year horizon, the effect cumulates to around 0.12 additional logins, as shown in Table 4. The magnitude of these estimates compares to a baseline value of the attention measure of around 6. The parallel trends in logins in the pre-period across households with different gains support our causal interpretation of these estimates.

These results resonate with existing empirical findings that stockholders reduce attention to investment accounts on days and weeks when markets are down (Sicherman et al., 2016). However, our results have a markedly different flavor by showing that financial attention is endogenous to stock market wealth shocks not just in the very short run, but over horizons as long as two years.

We conduct the same battery of robustness tests as for spending and net investment and report the results in Table A1 in the Appendix. The baseline results are generally highly robust to modifications of the sample and the specification irrespective of the horizon.

7 Conclusion

In this paper, we study how shocks to stock market wealth affect consumer spending over the short and the medium term. Compared to earlier papers, we break new ground by using granular data from a large retail bank to obtain precise measures of consumer spending and stock market portfolios at a high frequency and by developing a novel empirical framework where dynamic spending responses are identified from lottery-like variation in gains across households with ex ante similar portfolios.

We find that households adjust spending immediately and persistently in response to stock market gains. The magnitude of the responses implies that today’s stock market gains feed into

present and future household consumption at a relatively constant rate of approximately 4% per year. This is consistent with the permanent income hypothesis where households respond to windfalls by adjusting spending by a small amount in the present and all future periods.

We also find significant responses to stock market wealth shocks on two other behavioral margins. First, there is an immediate effect on investment decisions. Households convert paper gains into cash dividends through disinvestment, which may serve to finance increased spending and rebalance the portfolio across safe and risky assets. Second, there is a strong and lasting effect on financial attention: today's gains induce households to log on to their accounts more often in the present and future periods.

We find striking heterogeneity in these estimates when we split the sample by ex ante liquid wealth: both spending and investment responses are many times larger for the least liquid households than for the most liquid ones. The results suggest that many stockholders consume their gains relatively quickly; however, this has only minor implications for the real economy, as this group have small portfolios and therefore account for a modest share of aggregate stock market wealth.

We find much less heterogeneity when we split the sample by ex ante financial attention as even the least attentive households exhibit significant spending and disinvestment responses in the short run. This suggests that stockholders generally pay sufficient attention to stock prices to account for wealth shocks in their financial planning and that inattention does not play an important role in dampening household responses to stock market cycles.

References

- [1] Aladangady, A., 2017. "Housing wealth and consumption: Evidence from geographically-linked microdata." *American Economic Review* 107(11), p. 3415-46.
- [2] Alstadsæter, A., Johannesen, N. and Zucman, G., 2019. "Tax evasion and inequality." *American Economic Review* 109(6), p. 2073-2103.
- [3] Andersen, A.L., Jensen, A.S., Johannesen, N., Kreiner, C.T., Leth-Petersen, S. and Sheridan, A., 2020. "How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Data Sets." Working paper.
- [4] Andrei, D. and Hasler, M., 2015. "Investor attention and stock market volatility." *Review of Financial Studies* 28(1), p. 33-72.
- [5] Auclert, A., 2019. "Monetary policy and the redistribution channel." *American Economic Review* 109(6), p. 2333-67.
- [6] Auclert, A., Rognlie, M. and Straub, L., 2020. "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model" NBER working paper no. 26647.
- [7] Baker, S.R., 2018. "Debt and the response to household income shocks: Validation and application of linked financial account data." *Journal of Political Economy* 126(4), p. 1504-1557.
- [8] Bräuer, K., Hackethal, A. and Hanspal, T., 2020. "Consuming dividends." SAFE Working Paper.
- [9] Brunnermeier, M.K. and Parker, J.A., 2005. "Optimal expectations." *American Economic Review* 95(4), p. 1092-1118.
- [10] Calvet, L.E., Campbell, J.Y. and Sodini, P., 2007. "Down or out: Assessing the welfare costs of household investment mistakes." *Journal of Political Economy* 115(5), p. 707-747.
- [11] Calvet, L.E., Campbell, J.Y. and Sodini, P., 2009. "Fight or flight? Portfolio rebalancing by individual investors." *Quarterly Journal of Economics* 124(1), p. 301-348.
- [12] Campbell, J.Y. and Cochrane, J.H., 1999. "By force of habit: A consumption-based explanation of aggregate stock market behavior." *Journal of Political Economy* 107(2), p. 205-251.
- [13] Carroll, C.D., 1997. "Buffer-stock saving and the life cycle/permanent income hypothesis." *Quarterly Journal of Economics* 112(1), p. 1-55.
- [14] Carroll, C.D., Otsuka, M. and Slacalek, J., 2011. "How large are housing and financial wealth effects? A new approach." *Journal of Money, Credit and Banking* 43(1), p. 55-79.

- [15] Cesarini, D., Lindqvist, E., Östling, R. and Wallace, B., 2016. “Wealth, health, and child development: Evidence from administrative data on Swedish lottery players.” *Quarterly Journal of Economics* 131(2), p. 687-738.
- [16] Cesarini, D., Lindqvist, E., Notowidigdo, M.J. and Östling, R., 2017. “The effect of wealth on individual and household labor supply: evidence from Swedish lotteries.” *American Economic Review* 107(12), p. 3917-46.
- [17] Chetty, R. and Szeidl, A., 2016. “Consumption commitments and habit formation.” *Econometrica* 84(2), p. 855-890.
- [18] Chodorow-Reich, G., Nenov, P.T. and Simsek, A., 2021. “Stock market wealth and the real economy: A local labor market approach.” *American Economic Review* 111(5), p. 1613-57.
- [19] Cieslak, A. and Vissing-Jorgensen, A., 2020. “The Economics of the Fed Put.” NBER Working Paper (forthcoming in *Review of Financial Studies*).
- [20] Di Maggio, M., Kermani, A. and Majlesi, K., 2020. “Stock market returns and consumption.” *Journal of Finance* 75(6), p. 3175-3219.
- [21] Dynan, K.E. and Maki, D.M., 2001. “Does stock market wealth matter for consumption?” Working paper.
- [22] Fagereng, A., Gottlieb, C. and Guiso, L., 2017. “Asset market participation and portfolio choice over the life-cycle.” *Journal of Finance* 72(2), p. 705-750.
- [23] Fagereng, A., Guiso, L., Malacrino, D. and Pistaferri, L., 2020. “Heterogeneity and persistence in returns to wealth.” *Econometrica* 88(1), p. 115-170.
- [24] Fagereng, A., Holm, M.B. and Natvik, G.J., 2021. “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* (forthcoming).
- [25] Fama, E.F. and French, K.R., 2012. “Size, value, and momentum in international stock returns.” *Journal of Financial Economics* 105(3), p. 457-472.
- [26] Friedman, M., 1957. “The permanent income hypothesis”. In “A theory of the consumption function” (pp. 20-37). Princeton University Press.
- [27] Gabaix, X., 2019. “Behavioral inattention”. In “Handbook of Behavioral Economics: Applications and Foundations” 1 (Vol. 2, pp. 261-343). North-Holland.
- [28] Ganong, P. and Noel, P., 2019. “Consumer spending during unemployment: Positive and normative implications.” *American Economic Review* 109(7), p. 2383-2424.
- [29] Ganong, P., Jones, D., Noel, P., Greig, F., Farrell, D. and Wheat, C., 2020. “Wealth, race, and consumption smoothing of typical income shocks.” NBER Working Paper no. 27552.
- [30] Gargano, A. and Rossi, A.G., 2018. “Does it pay to pay attention?” *Review of Financial Studies* 31(12), p. 4595-4649.

- [31] Gelman, M., Kariv, S., Shapiro, M.D., Silverman, D. and Tadelis, S., 2014. “Harnessing naturally occurring data to measure the response of spending to income.” *Science* 345(6193), p. 212-215.
- [32] Gomes, F. and Michaelides, A., 2005. “Optimal life-cycle asset allocation: Understanding the empirical evidence.” *Journal of Finance* 60(2), p. 869-904.
- [33] Harvey, C.R. and Siddique, A., 2000. “Conditional skewness in asset pricing tests.” *Journal of Finance* 55(3), p. 1263-1295.
- [34] Jordà, Ò., 2005. “Estimation and inference of impulse responses by local projections.” *American Economic Review* 95(1), p. 161-182.
- [35] Johnson, D.S., Parker, J.A. and Souleles, N.S., 2006. “Household expenditure and the income tax rebates of 2001.” *American Economic Review* 96(5), p. 1589-1610.
- [36] Kaplan, G. and Violante, G.L., 2014. “A model of the consumption response to fiscal stimulus payments.” *Econometrica* 82(4), p. 1199-1239.
- [37] Kaplan, G. and Violante, G.L., 2018. “Microeconomic heterogeneity and macroeconomic shocks.” *Journal of Economic Perspectives* 32(3), p. 167-94.
- [38] Kaplan, G., Moll, B. and Violante, G.L., 2018. “Monetary policy according to HANK.” *American Economic Review* 108(3), p. 697-743.
- [39] Karlsson, N., Loewenstein, G. and Seppi, D., 2009. “The ostrich effect: Selective attention to information.” *Journal of Risk and Uncertainty* 38(2), p. 95-115.
- [40] Kleven, H., Knudsen, M., Kreiner, C., Pedersen, S. and Saez, E., 2011. “Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark.” *Econometrica* 79, p. 651-692.
- [41] Kueng, L., 2018. “Excess sensitivity of high-income consumers.” *Quarterly Journal of Economics* 133(4), p. 1693-1751.
- [42] Lettau, M. and Ludvigson, S.C., 2004. “Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption.” *American Economic Review* 94(1), p. 276-299.
- [43] Mankiw, N.G. and Zeldes, S.P., 1991. “The consumption of stockholders and nonstockholders.” *Journal of Financial Economics* 29(1), p. 97-112.
- [44] Meyer, S. and Pagel, M., 2019. “Fully closed: Individual responses to realized gains”. NBER Working Paper No. 25542 (forthcoming in *Journal of Finance*).
- [45] Mian, A., Rao, K. and Sufi, A., 2013. “Household balance sheets, consumption, and the economic slump.” *Quarterly Journal of Economics* 128(4), p. 1687-1726.

- [46] Olafsson, A. and Pagel, M., 2017. “The ostrich in us: Selective attention to financial accounts, income, spending, and liquidity.” NBER Working paper No. 23945.
- [47] Paiella, M. and Pistaferri, L., 2017. “Decomposing the wealth effect on consumption.” *Review of Economics and Statistics* 99(4), p. 710-721.
- [48] Poterba, J.M., 2000. “Stock market wealth and consumption.” *Journal of Economic Perspectives* 14(2), p. 99-118.
- [49] Reis, R., 2006. “Inattentive consumers.” *Journal of Monetary Economics* 53(8), p. 1761-1800.
- [50] Sheridan, A., Andersen, A. L., Hansen, E. T. and Johannesen, N., 2020. “Social distancing laws cause only small losses of economic activity during the COVID-19 pandemic in Scandinavia.” *Proceedings of the National Academy of Sciences* 117(34), p. 20468-20473.
- [51] Sicherman, N., Loewenstein, G., Seppi, D.J. and Utkus, S.P., 2016. “Financial attention.” *Review of Financial Studies* 29(4), p. 863-897.
- [52] Zeldes, S.P., 1989. “Consumption and liquidity constraints: an empirical investigation.” *Journal of Political Economy* 97(2), p. 305-346.

Table 1: Summary statistics. The table shows summary statistics for our estimation sample. Starting from the sample of all households with accounts at Danske Bank, we arrive at the estimation sample by excluding observations where (i) the value of the portfolio is below 100 kroner; (ii) the number of monthly spending transactions is below five; (iii) the portfolio includes an investment in a firm that is also the primary employer of one household member; (iv) the only stock market investment is shares in Danske Bank. Portfolio size is the dollar value of the portfolio. Portfolio risk is the standard deviation of monthly portfolio returns over the past 12 months. Expected return is the mean of monthly portfolio returns over the past 12 months.

	Mean	SD
Demographics		
Age (years)	60.7	17.5
Highest education:		
- High school	0.53	0.50
- College	0.28	0.45
- Graduate	0.17	0.38
Has partner	0.54	0.50
Is exclusive DB customer	0.49	0.50
Characteristics of stock market portfolio		
Portfolio size	80,182	156,906
Portfolio risk	0.047	0.052
Expected return	0.008	0.022
Number of securities	4.2	4.2
Share of financial wealth	0.48	0.33
Gains, losses and investment (monthly in \$)		
Actual gains and losses	587	4,809
Passive gains and losses	594	4,817
Net investment	47	3,304
Dividends	180	1,286
Income, income and liquidity (monthly in \$)		
Total income	6,456	4,790
Labor income	3,806	5,458
Spending	3,598	2,768
Deposits	49,930	64,373
Portfolio links to labor market		
Share of portfolio in own industry	0.04	0.17
Has equity investments in own industry	0.08	0.28
Has unlisted assets in portfolio	0.10	0.29
Number of households	390,151	-
Number of household-month observations	13,582,087	-

Table 2: Cumulative marginal propensity to consume. The table shows the estimated marginal propensity to consume out of stock market gains cumulated over different horizons (columns) and for different empirical specifications (rows). *Full sample* means that households with investments in a firm that is also the main employer of a household member are allowed to enter the estimation sample. *Exclusive* means that households with accounts in other banks than Danske Bank are excluded from the estimation sample. *Excl. own industry* means that we ignore investments in industries where household members are employed when we construct the stock market portfolio. *No controls* means that we estimate the model without household-level controls. *Excl. unlisted* means that we drop households with unlisted shares. Standard errors are clustered at the household level.

	3 months	6 months	12 months	24 months	36 months
Main results					
Baseline	.0056 (.0018)	.016 (.0037)	.038 (.0079)	.0754 (.0197)	.1192 (.0373)
Robustness					
Full sample	.0066 (.0017)	.0187 (.0035)	.0427 (.0074)	.0791 (.0185)	.115 (.035)
Exclusive	.0088 (.0037)	.022 (.0073)	.0541 (.0147)	.1053 (.0318)	.1665 (.0589)
Excl. own industry	.0053 (.0019)	.0153 (.004)	.0393 (.0085)	.0873 (.0213)	.1159 (.0405)
No controls	.0057 (.0018)	.0162 (.0037)	.0381 (.0079)	.0768 (.0197)	.1206 (.0373)
Excl. unlisted	.0053 (.0019)	.0147 (.004)	.0388 (.0084)	.0852 (.0209)	.1229 (.0392)

Table 3: Heterogeneity in liquidity, age, returns and attention. The table shows the mean of liquidity, household age, absolute return and financial attention by tercile of each of these variables. For each household-month observation, we measure liquidity as the sum of deposits, stocks and other securities held at the beginning of the month relative to average monthly income; household age as the age of the oldest household member at the beginning of the month; absolute return as the absolute size of the stock market portfolio return during the month; and financial attention as the number of days that a member of the household logged into their Danske Bank accounts in an average month. In each dimension, we rank the observations within each month and split them into three equally sized groups: the first, second and third terciles. The reported means are taken within these groups across all months in the sample. Standard errors are clustered at the household level.

	First tercile	Second tercile	Third tercile
Liquidity	3.2	14.5	75.0
Household age	40	64	79
Absolute return	0.6%	1.8%	7.7%
Financial attention	0.9	5.0	12.1

Table 4: Cumulative investment and attention responses. The table shows the estimated responses to a \$1 stock market gain in the form of net investment and account logins cumulated over different horizons (columns). Standard errors are clustered at the household level.

	3 months	6 months	12 months	24 months
Other outcomes				
Net investment	-.0419 (.0096)	-.1191 (.0207)	-.1448 (.0431)	-.2057 (.0943)
Logins	.0372 (.0015)	.0631 (.0033)	.1169 (.0076)	.1794 (.0224)

Figure 1: Data validation. The upper panel shows a bin scatter of the value of total portfolios as reported by financial institutions to the Danish tax authorities (x-axis) and the value of portfolios held at Danske Bank as observed in the customer dataset (y-axis). The lower panel shows a bin scatter of annual income as reported on the tax return (x-axis) and annual spending as measured in the customer dataset from Danske Bank (y-axis). Both figures use the estimation sample defined in Section 2.3.

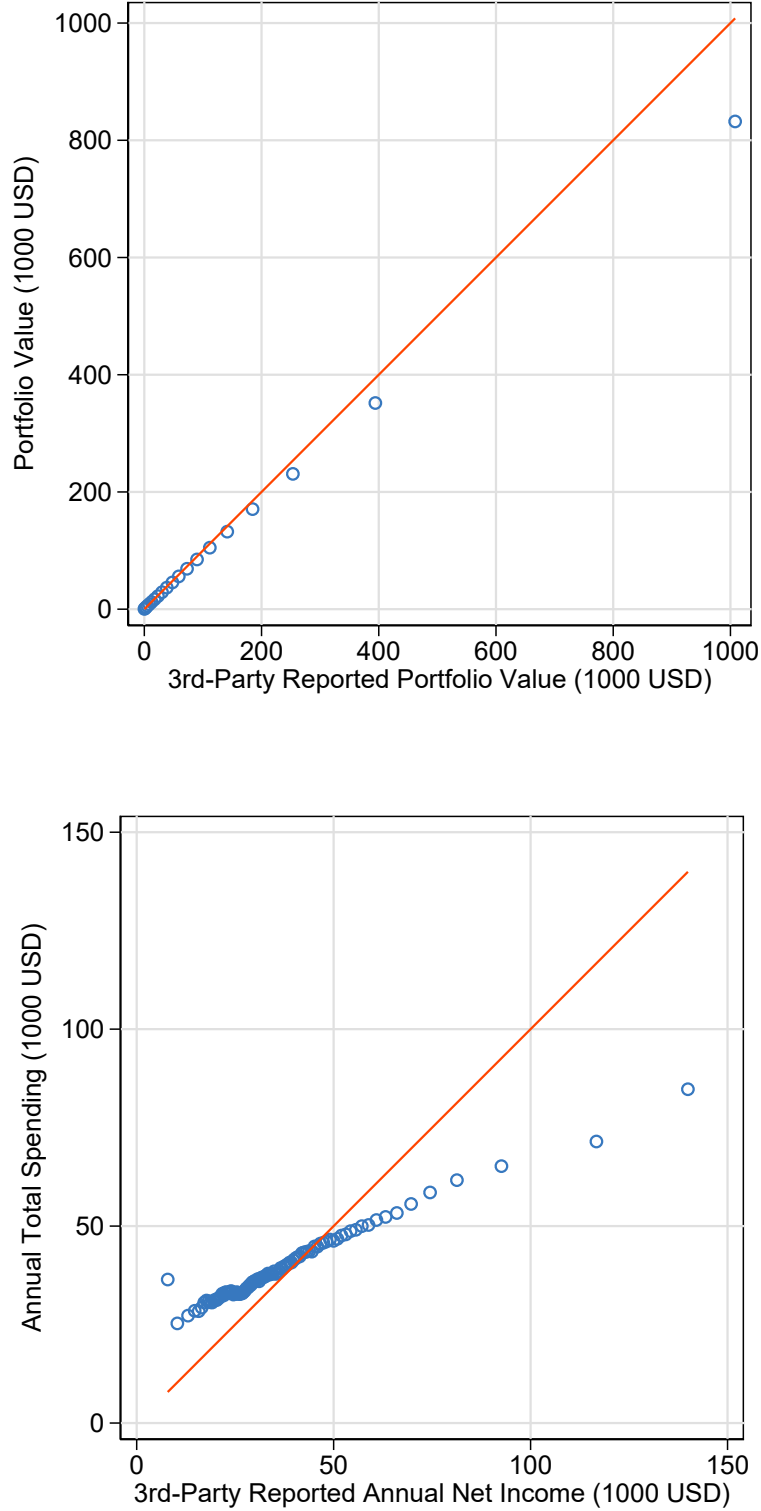


Figure 2: Variation in returns. The figure shows monthly returns for the leading Danish stock market index OMXC20 (black line), the average monthly portfolio return within our analysis sample (orange line) and the 5th and 95th percentile of monthly portfolio returns in our analysis sample (shaded area).

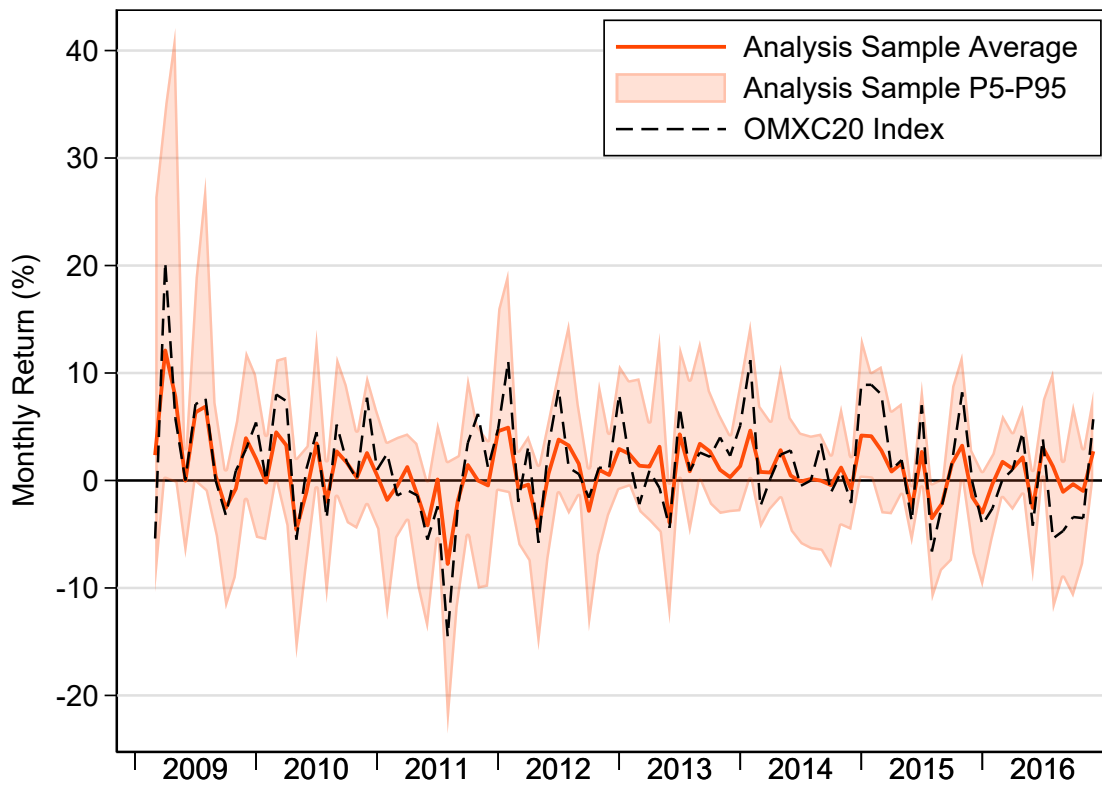


Figure 3: Lottery tests. The figure shows binned scatterplots of five different observables against passive gains unconditionally (blue circles) and conditional on controls for portfolio characteristics: portfolio size, risk and expected return (orange diamonds). The five observables are: (a) age, (b) deposits, (c) monthly total income, (d) share of risky assets in the portfolio and (e) monthly total spending. The conditional scatterplots have added the sample mean back in so that levels are comparable to the unconditional scatterplots.

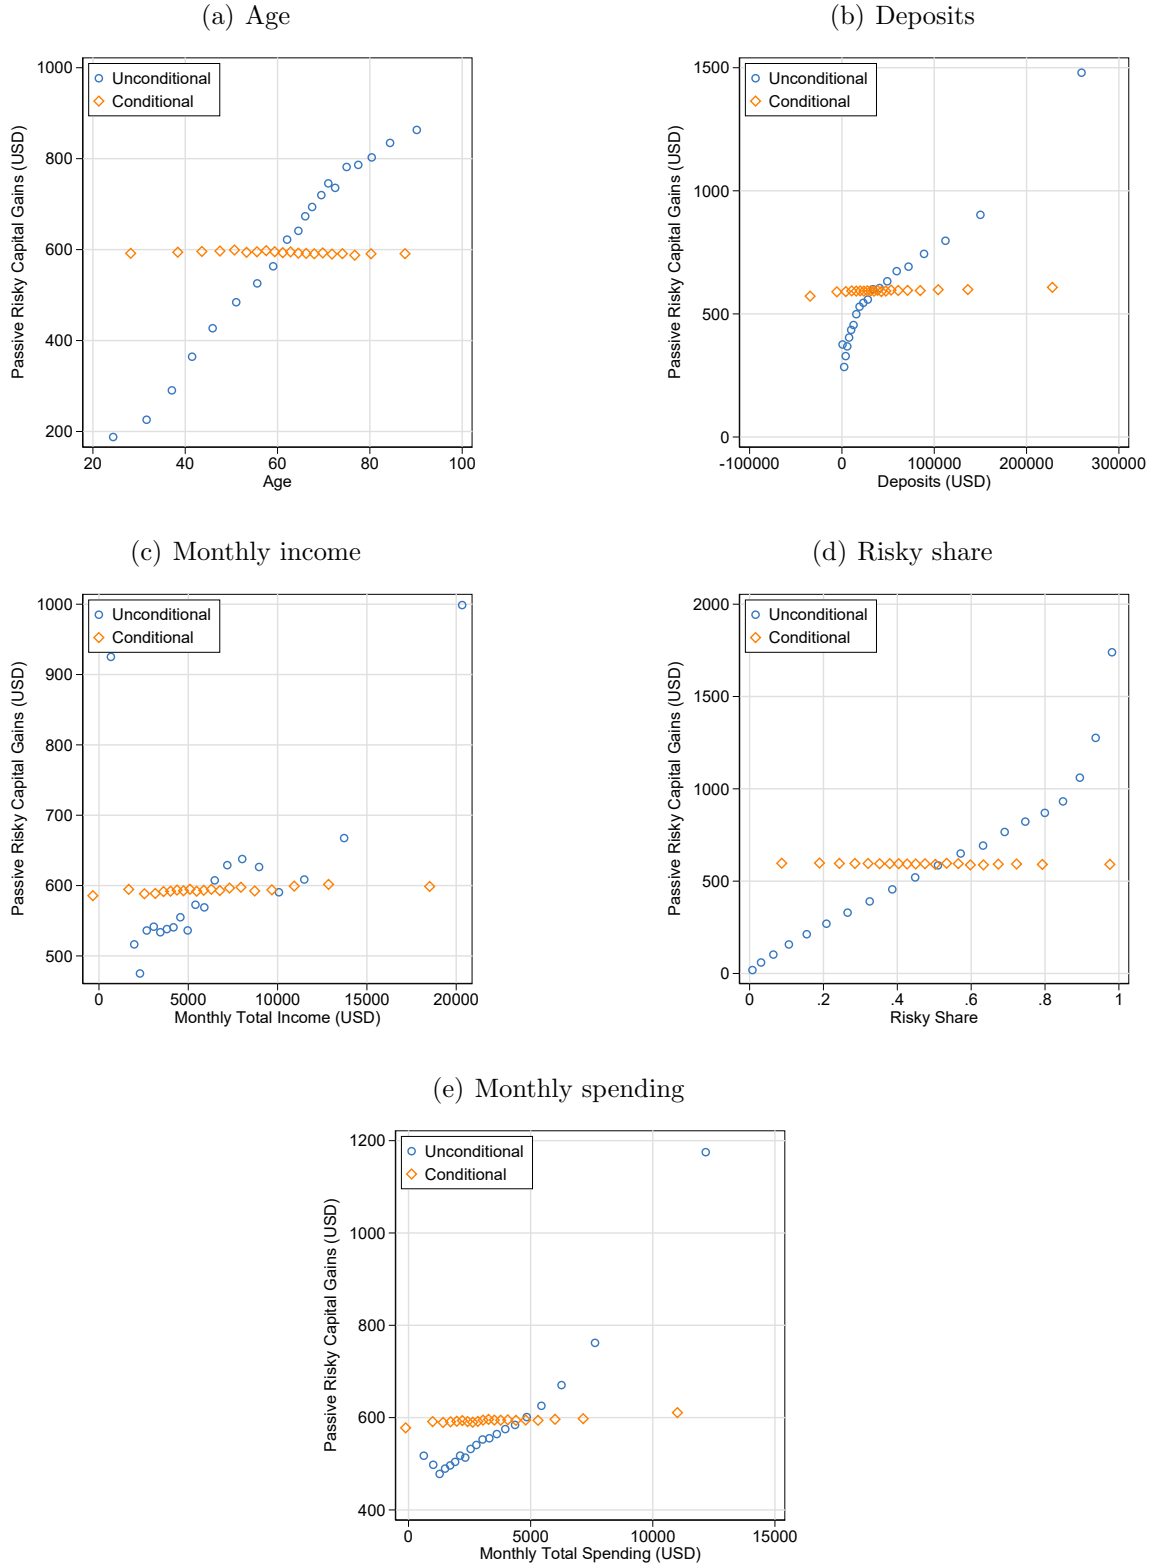


Figure 4: Passive portfolio values. The figure shows the estimated change in the value of the portfolio held at the beginning of period $t + 1$ caused by a \$1 passive gain in period $t + 1$. We obtain the estimates from equation (3) using the passive change in the value of the portfolio held at the beginning of month $t + 1$ as the dependent variable. The confidence intervals are based on standard errors clustered at the household level.

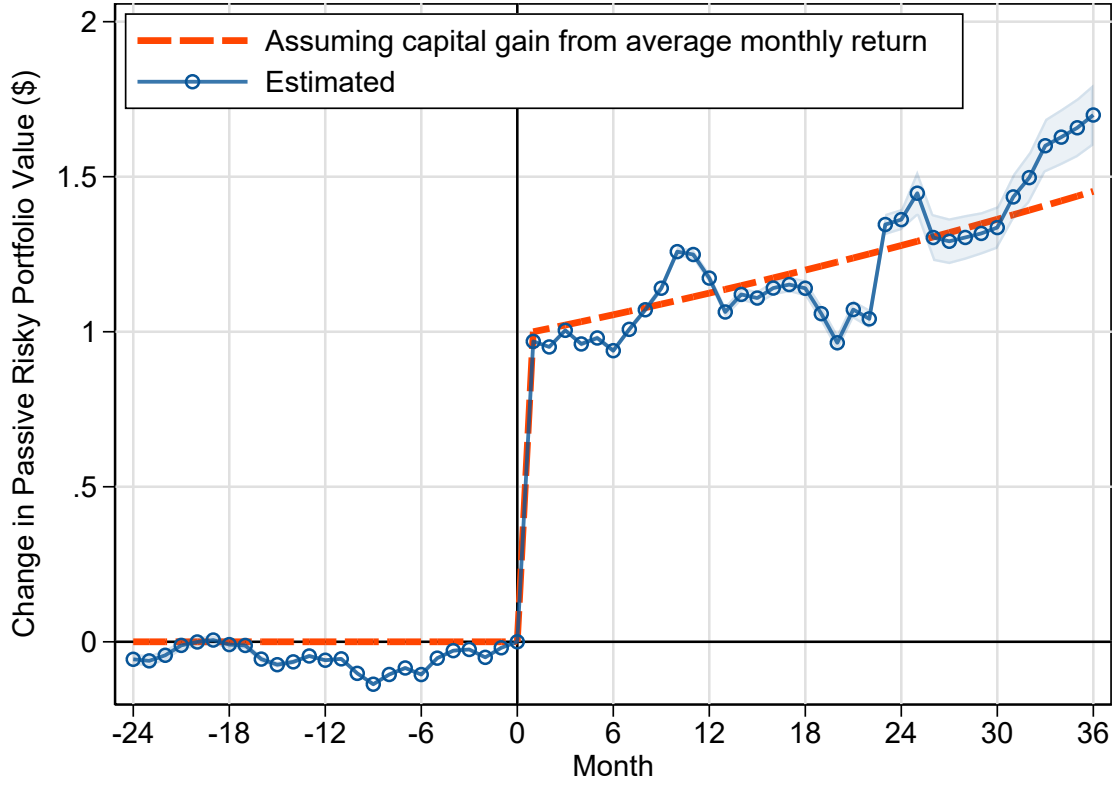


Figure 5: Marginal propensity to consume. The figure illustrates our dynamic estimates of the marginal propensity to consume: the change in spending between month t and $t + h$ caused by a \$1 stock market gain in month $t + 1$. On the x-axis is the time horizon h . On the y-axis is the marginal propensity to consume. The confidence intervals are based on standard errors clustered at the household level.

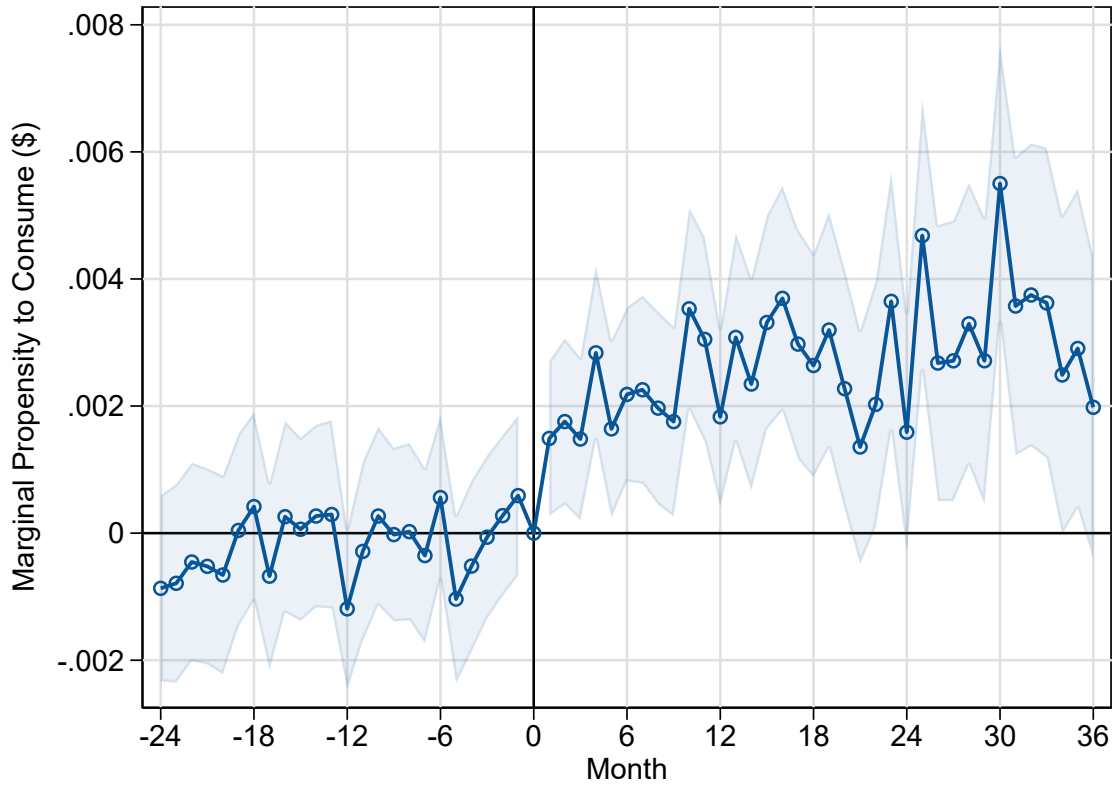


Figure 6: Robustness of baseline estimates The figure illustrates our dynamic estimates of the marginal propensity to consume for different variations of the sample and the specification. *Full sample* means that households with investments in a firm that is also the main employer of a household member are allowed to enter the estimation sample. *Exclusive* means that households with accounts in other banks than Danske Bank are excluded from the estimation sample. *Excl. own industry* means that we ignore investments in industries where household members are employed when we construct the stock market portfolio. *No controls* means that we estimate the model without household-level controls. *Excl. unlisted* means that we drop households with unlisted shares. The confidence intervals are based on standard errors clustered at the household level.

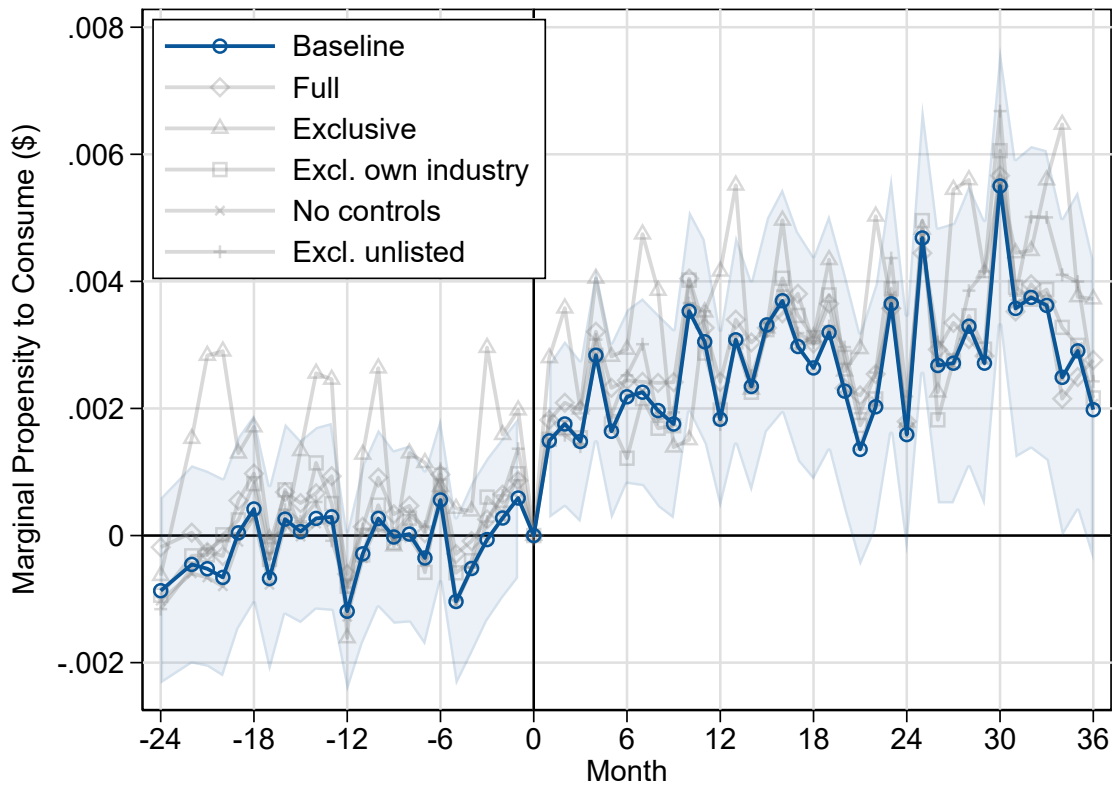
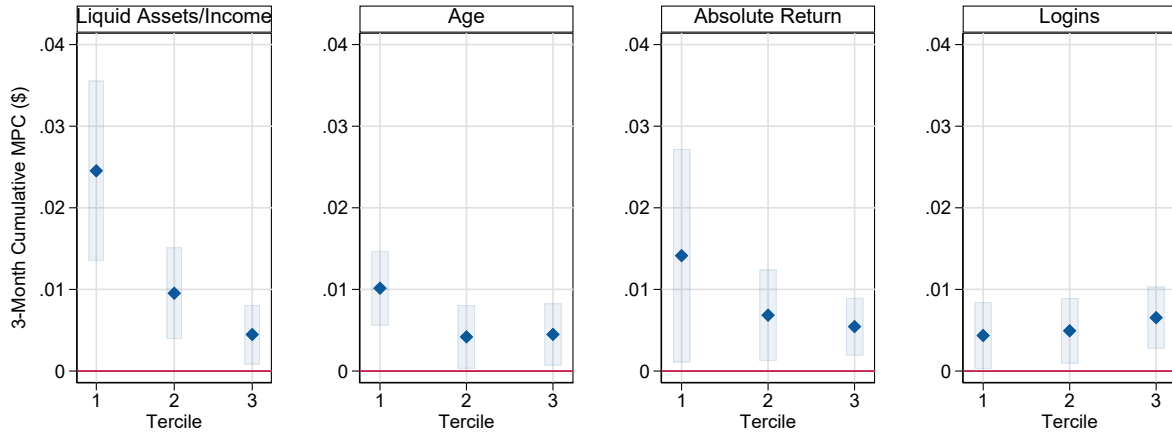


Figure 7: Heterogeneity in the marginal propensity to consume. The figure shows heterogeneity in the estimated marginal propensity to consume out of stock market gains cumulated over a three-month horizon (top panel) and one-year and two-year horizons (bottom panel). The dimensions of heterogeneity are: (i) liquid assets as a fraction of income, (ii) household age, (iii) the absolute size of the portfolio return and (iv) the frequency of days with account logins. The results are obtained from a modified version of the baseline model where gains are interacted with three indicators capturing the heterogeneity in the dimension of interest. The confidence intervals are based on standard errors clustered at the household level.

(a) Short-run responses



(b) Medium-run responses

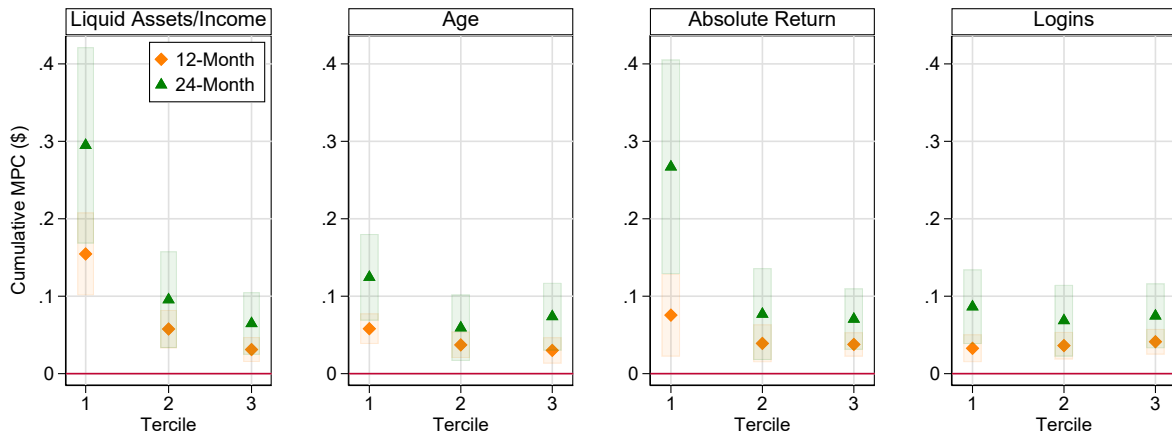


Figure 8: Dynamic investment and attention responses. The figure illustrates our dynamic estimates of the investment and attention responses: the increase in net investment (upper panel) and account logins (lower panel) between month t and $t + h$ caused by a \$1 stock market gain in month $t + 1$. On the x-axis is the time horizon h . On the y-axis is the marginal propensity to consume. The confidence intervals are based on standard errors clustered at the household level.

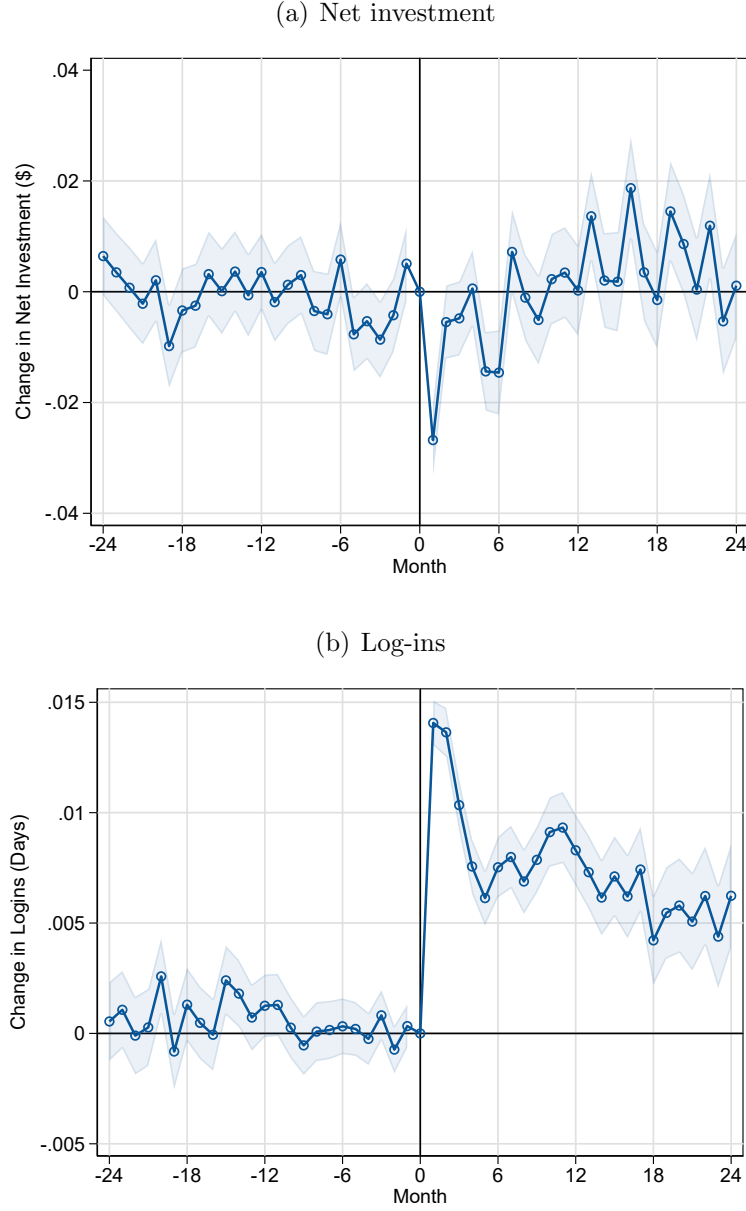
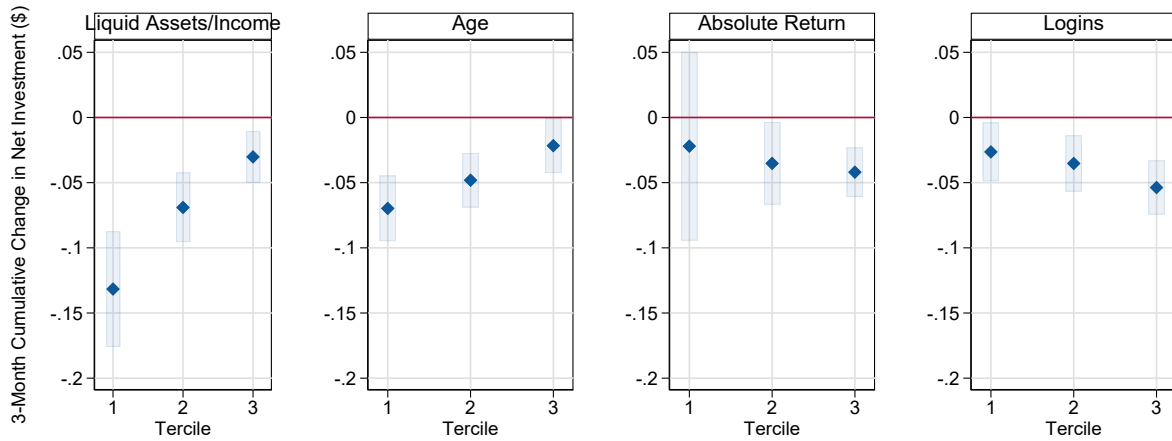
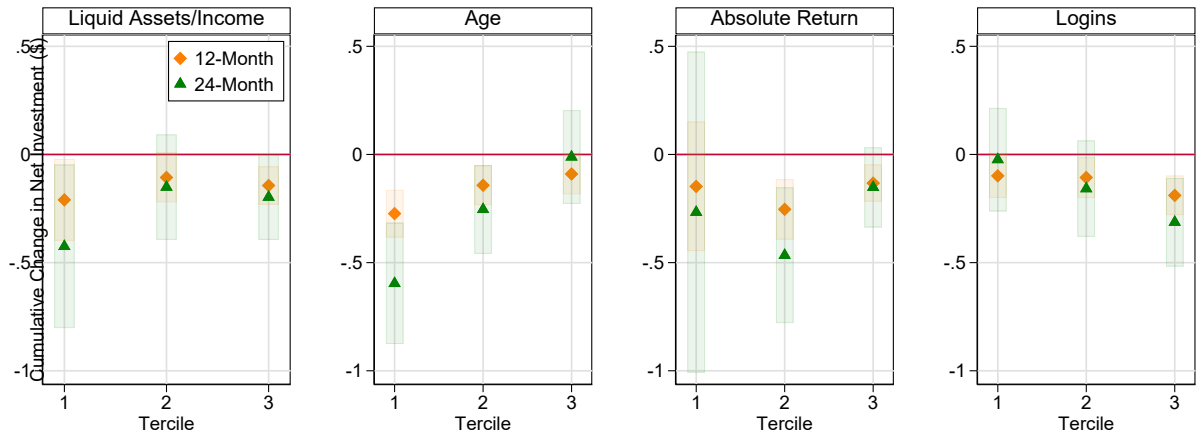


Figure 9: Heterogeneity in investment responses. The figure illustrates the heterogeneity in the estimated investment responses to a \$1 stock market gain cumulated over a three-month horizon (top panel) and one-year and two-year horizons (bottom panel). The dimensions of heterogeneity are: (i) liquid assets as a fraction of income, (ii) household age, (iii) the absolute size of the portfolio return and (iv) the frequency of days with account logins. The results are obtained from a modified version of the baseline model where gains are interacted with three indicators capturing the heterogeneity in the dimension of interest. The confidence intervals are based on standard errors clustered at the household level.

(a) Short-run responses



(b) Medium-run responses



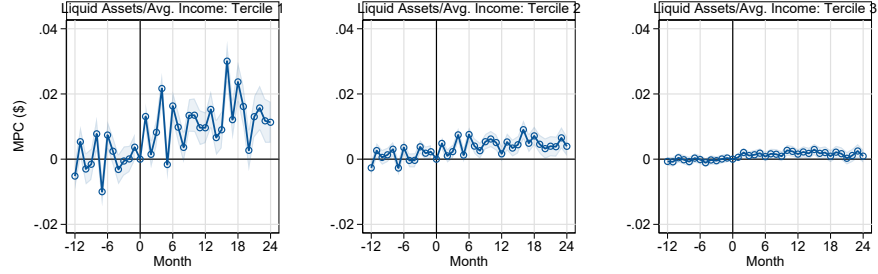
APPENDIX

Table A1: Robustness. The table shows the estimated responses in terms of net investment (upper panel) and account logins (bottom panel) cumulated over different horizons (columns) and for different empirical specifications (rows). *Full sample* means that households with investments in a firm that is also the main employer of a household member are allowed to enter the estimation sample. *Exclusive* means that households with accounts in other banks than Danske Bank are excluded from the estimation sample. *Excl. own industry* means that we ignore investments in industries where household members are employed when we construct the stock market portfolio. *No controls* means that we estimate the model without household-level controls. *Excl. unlisted* means that we drop households with unlisted shares.

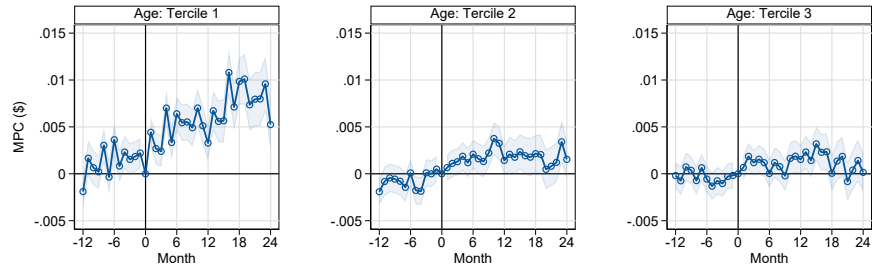
	3 months	6 months	12 months	24 months
Net investment				
- Baseline	-.0419 (.0096)	-.1191 (.0207)	-.1448 (.0431)	-.2057 (.0943)
- Full sample	-.0409 (.0091)	-.1085 (.0195)	-.1292 (.0402)	-.152 (.0872)
- Exclusive	-.0533 (.0196)	-.1111 (.038)	-.2069 (.0763)	-.4278 (.1545)
- Excl. own industry	-.042 (.0106)	-.1355 (.0226)	-.1913 (.047)	-.2947 (.1027)
- No controls	-.042 (.0096)	-.1184 (.0207)	-.1428 (.0432)	-.1999 (.0943)
- Excl. unlisted	-.0385 (.0105)	-.1182 (.0226)	-.1635 (.0461)	-.245 (.1014)
Logins				
- Baseline	.0372 (.0015)	.0631 (.0033)	.1169 (.0076)	.1794 (.0224)
- Full sample	.0391 (.0014)	.0656 (.0031)	.1155 (.0072)	.1695 (.021)
- Exclusive	.0408 (.0033)	.0644 (.0067)	.0995 (.0148)	.1664 (.0362)
- Excl. own industry	.034 (.0016)	.0568 (.0036)	.1076 (.0082)	.1634 (.0243)
- No controls	.0372 (.0015)	.063 (.0033)	.1163 (.0076)	.1778 (.0224)
- Excl. unlisted	.0356 (.0016)	.0619 (.0035)	.1145 (.0082)	.1981 (.0241)

Figure A1: Heterogeneity in dynamic spending responses. The figure illustrates the heterogeneity in our dynamic estimates of the marginal propensity to consume: the increase in spending between month t and $t + h$ caused by a \$1 stock market gain in month $t + 1$. Standard errors are clustered at the household level.

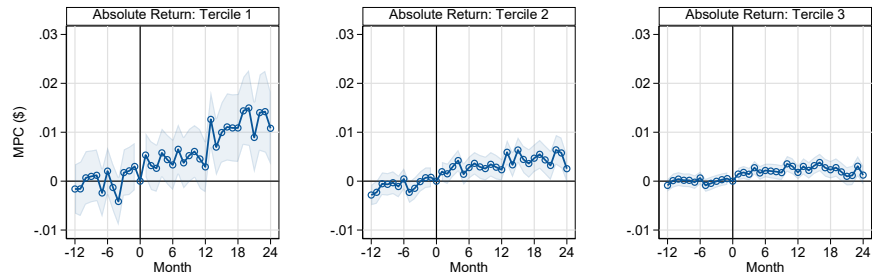
(a) Liquidity



(b) Age



(c) Absolute return



(d) Financial attention

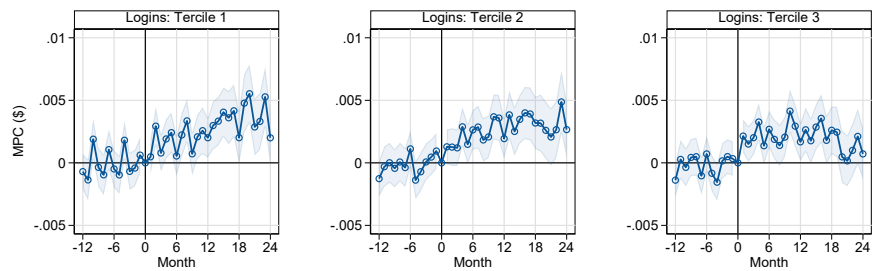
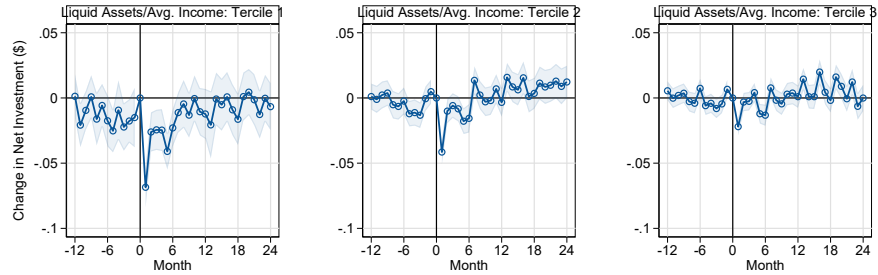
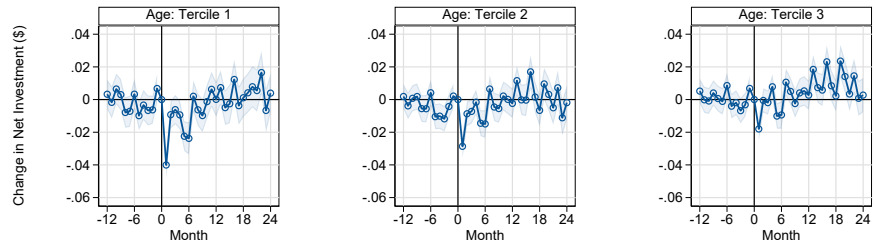


Figure A2: Heterogeneity in dynamic investment responses. The figure illustrates the heterogeneity in our dynamic estimates of the investment responses: the change in net investment between month t and $t + h$ caused by a \$1 stock market gain in month $t + 1$. Standard errors are clustered at the household level.

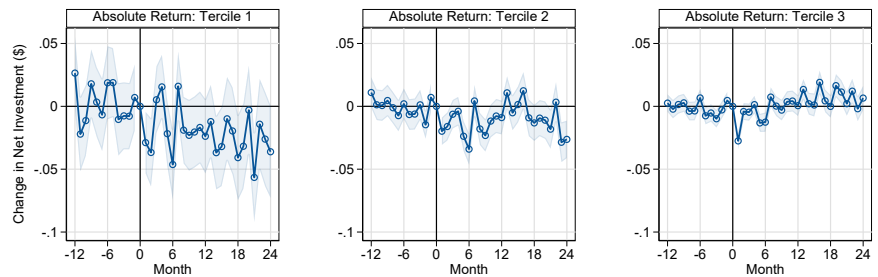
(a) Liquidity



(b) Age



(c) Absolute return



(d) Financial attention

