AER: Insights 2024, 6(3): 434–452
https://doi.org/10.1257/aeri.20230382

Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi Lotteries on the Stock Market†

By Asger Lau Andersen, Niels Johannesen, and Adam Sheridan*

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 monthly responses cumulate to marginal propensities to consume of 4.4 percent over one year and 16 percent over three years. The results suggest that inattention attenuates household responses to stock market cycles over horizons as long as one year. (*JEL* D12, E21, G11, G14, G51)

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 US 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

*Andersen: Center for Economic Behavior and Inequality, University of Copenhagen (email: asger.lau.andersen@econ.ku.dk); Johannesen: Centre for Business Taxation, Oxford University and the Center for Economic Behavior and Inequality (CEBI), University of Copenhagen (email: niels.johannesen@sbs.ox.ac.uk); Sheridan: Amazon (email: afmsheridan@gmail.com). Peter Klenow was the coeditor for this article. We are extremely grateful to key employees at Danske Bank for their help. All individual data used in this analysis have 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. Sheridan worked on this project prior to joining Amazon. We gratefully acknowledge financial support from the Danish National Research Foundation, the Carlsberg Foundation (grant CF16-0563), the Independent Research Fund Denmark (grant DFF-4182-00101), the Economic Policy Research Network, and Innovation Fund Denmark. The activities of CEBI are financed by the Danish National Research Foundation (grant DNRF134).

†Go to https://doi.org/10.1257/aeri.20230382 to visit the article page for additional materials and author disclosure statement(s).

policymakers deciding how strongly to lean against the stock market to stabilize aggregate demand (Cieslak and Vissing-Jorgensen 2021) and for macroeconomists 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), buffer stock models imply a larger but temporary increase (Carroll 1997), and theories of behavioral inattention suggest a delayed response (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 350,000 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 (Gelman et al. 2014). 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 contain information about account log-ins, which allows us to explore the role of financial attention (Sicherman et al. 2016). Finally, we add data from government registers to ascertain the completeness and representativeness of the bank data (Baker 2018).

Our empirical strategy exploits the comprehensive data on stock market portfolios to isolate lottery-like variation in gains. The main idea is to compare households that 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 portfolio owner characteristics.

Drawing on this idea, we develop a model that identifies the dynamic spending response to stock market wealth shocks. The explanatory variable is the household's stock market gains in the current month. The dependent variable is the change in the household's spending from the previous to the current month or, to study lagged responses, the change in spending over longer horizons. We augment this framework in two ways. First, we instrument *actual* gains, which are endogenous to trading during the month, with *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: portfolio value (size), volatility of past returns (risk), and level of past returns (expected return). Thus, identification requires that the gains of the portfolios held at the beginning of the month are randomly assigned conditional on the portfolios' size, risk, and expected return.

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 in the first month, around 4.4 cents cumulated over one year, and around 16.4 cents cumulated over three years. This is broadly 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. However, rather than a simple level shift in monthly spending, our estimates indicate a somewhat gradual adjustment to a higher spending level.

We probe the validity and robustness of these results in several ways. First, we estimate the model for negative horizons and find no correlation between current gains and *past* changes in spending. The parallel spending trajectories in the preperiod support our causal interpretation of the diverging spending trajectories in the postperiod. Second, we address specific identification concerns, for example, about spending and investments through other banks and about confounding shocks to expected future wages.

Finally, we investigate how the spending responses vary with household characteristics in key dimensions. The heterogeneity by liquidity is most striking: spending responses are persistently around four times larger for the least liquid tercile of the households than for the most liquid tercile. However, we also find notable heterogeneity by financial attention: spending responses are only half as large for the least attentive tercile of the households as for the most attentive tercile over the first three months, and the difference, while diminishing over time, remains significant over a one-year horizon. The latter result is suggestive that inattention may play a role in explaining the gradual spending adjustment to stock market gains.

Our first contribution is methodological. Our research design identifies the wealth effect on consumption by comparing households that hold portfolios with similar characteristics but earn different returns due to random factors. Most existing microstudies in the literature use much cruder data from household surveys, in which, notably, gains are poorly measured and that identify wealth effects by comparing households whose portfolios differ in terms of size and assume that they earn the same return (e.g., Dynan and Maki 2001; Paiella and Pistaferri 2017). Di Maggio, Kermani, and Majlesi (2020) use administrative tax data where consumption can be approximated from income and balance sheets. They estimate the same-year effect of stock market gains on consumption but do not characterize the dynamics of the adjustment—whether it is immediate or sluggish, persistent or transient—which is a key focus of our paper. 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.[1]

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 4.4 percent of their aggregate stock market gains over one year is somewhat larger

[1] Meyer and Pagel (2019) and Bräuer, Hackethal, and Hanspal (2022) investigate how spending responds to forced realizations and dividend payouts but do not consider wealth shocks. Several studies use macrodata on stock markets and consumption (e.g., Lettau and Ludvigson 2004; Carroll, Otsuka, and Slacalek 2011).

than the estimate of Chodorow-Reich, Nenov, and Simsek (2021) and the central scenario in Poterba (2000), both of which are around 3 percent.[2]

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). While our results are consistent with a role for inattention, it is also noteworthy that we find immediate and persistent spending responses across all subsamples, including the most liquid and the least attentive ones.

Relatedly, we contribute to the literature on heterogeneity in the marginal propensity to consume, which is central to 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 we study stock owners who are relatively liquid. It resonates with existing work on stock market wealth shocks (Di Maggio, Kermani, and Majlesi 2020) and studies of spending responses to labor income shocks (Ganong et al. 2020), stimulus payments (Johnson, Parker, and Souleles 2006), unemployment shocks (Andersen et al. 2023), and lottery gains (Fagereng, Holm, and Natvik 2021).

We describe the data in Section I, develop the empirical framework in Section II, present the results in Section III, and conclude in Section IV.

I. Data

A. *Bank Data*

Our primary data source is the complete set of customer records of all personal customers at Danske Bank, the largest retail bank in Denmark, for the period 2009–2016 (Danske Bank 2018).

First, we use transaction-level data to construct a monthly measure of *consumer spending*, $C_{i,t}$, which aggregates purchases made with debit and credit cards, bill payments, in-store mobile payments, and cash withdrawals. We exclude transactions that are not associated with consumption, such as debt service and tax payments. Moreover, to ensure comparability across homeowners and renters, we exclude transactions related to housing expenditure, such as rent payments. The online Appendix provides more details and documents that our transaction-based approach yields spending levels that are similar to consumption surveys (see Figure A1).

Second, we use daily data on assets and security prices to construct a monthly measure of stock market gains, $G_{i,t}$. We start by delimiting each household's stock market portfolio. Following Calvet, Campbell, and Sodini (2007, 2009), we include direct equity investments and mutual fund shares but exclude direct bond investments and deposits. We include all liquid securities accounts but exclude retirement savings

[2] Other MPC estimates for stock market wealth vary considerably, from around zero (Paiella and Pistaferri 2017) to 10–15 percent (Dynan and Maki 2001). Di Maggio, Kermani, and Majlesi (2020) report no estimate for the full sample. Relatedly, recent MPC estimates for housing wealth shocks are 5–7 percent (Mian, Rao, and Sufi 2013) and 5 percent (Aladangady 2017).

accounts where assets typically cannot be liquidated without significant tax penalties.[3] Letting $a_{i,j,d}$ denote the number of security j held by household i on day d and letting $\Delta p_{j,d}$ denote the change in the price of security j on day d adjusted for stock splits, mergers of assets, dividend payouts, and other corporate events, we define stock market gains in month t as

$$G_{i,t} = \sum_{d \in t} \sum_{j} \Delta p_{j,d} \cdot a_{i,j,d}. \tag{1}$$

We note that, since $\Delta p_{j,d}$ 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.[4]

Third, we use the data on assets and security prices to measure the ex ante portfolio characteristics that serve as controls. Specifically, 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 12 months; and *expected returns* as the mean of this portfolio's returns over the past 12 months where the portfolio return is itself defined as gains in a given month measured relative to the portfolio size.

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. Several recent studies have used account log-ins to measure households' attention to private finances (Karlsson, Loewenstein, and Seppi 2009; Sicherman et al. 2016; Olafsson and Pagel 2017; Gargano and Rossi 2018).

B. *Government Register Data*

We merge the customer data from Danske Bank with administrative data from various government registers, which enables us to address a range of concerns about identification. The two data sources are merged on personal identification numbers.

From the population register (Statistics Denmark 2018a), we add household identifiers that link cohabiting couples and dependents. This allows us to aggregate all variables across members of the same household and estimate the marginal propensity to consume at the household level. If 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.

From tax registers (Statistics Denmark 2021a, b), we add end-of-year information about securities, deposits, and loans in all Danish financial institutions, collected by the government for tax enforcement purposes (Alstadsæter, Johannesen, and Zucman 2019). We use this information to identify households that have no accounts in banks other than Danske Bank. In robustness tests, we reestimate the model for

[3] Disregarding illiquid pension accounts may introduce a bias if portfolio choices in pension and nonpension accounts are correlated. We find that the potential bias is modest. A \$1 gain on nonpension accounts is associated with a \$0.05 gain on pension accounts (conditional on portfolio controls). Hence, even if illiquid gains on pension accounts feed into spending at the same rate as liquid gains on nonpension accounts, our MPC estimates are only overestimated by a factor of 1.05.

[4] Some related papers distinguish between capital gains and dividends (Di Maggio, Kermani, and Majlesi 2020).

this subsample of exclusive customers, which is less representative than the baseline sample but where secondary bank relations are not a potential confounder.

From employment registers (Statistics Denmark 2018b, c), we add links between individuals and their employers. Using data from Bureau van Dijk (2018), we create a crosswalk between employer identifiers in the employment register and security identifiers in the portfolio data and use it to identify households that 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 employee-shareholders may be affected by shocks to firm profitability not just through stock market gains but also through changes in expected future wages.

Finally, we construct variables that serve as household-level controls based on information from various administrative registers. From the education register (Statistics Denmark 2018f), we construct categorical variables indicating the highest level of education in the household: primary school, high school, college, or graduate degree.[5] From the population register, we construct a variable indicating the age of the oldest household member and a measure of household size. From income and employment registers (Statistics Denmark 2018b, d, e), we obtain data on monthly and annual income as well as the value of assets held at the end of each calendar year. This information is collected by the government for tax purposes.

C. *Estimation Sample*

To obtain our main estimation sample, we start from all households with securities accounts at Danske Bank. These households account for around 40 percent of the aggregate spending of the bank's customers. We then impose some mild sample restrictions. First, we exclude households whose portfolio values are negligible—that is, below 100 kroner (around US$15). Second, to ensure that a household is an active Danske Bank customer in a given month, we require that its adult members each have at least one spending transaction in that month, and at least five jointly. Third, we drop households that invest in a firm that is also the primary employer of a household member because of the potentially confounding effect of expected future wages. Finally, we drop households whose only investment is shares in Danske Bank. With these sample restrictions, we arrive at our baseline estimation sample with around 350,000 distinct households and around 13.5 million household-month observations.

The completeness of the bank data and the representativeness of the bank sample are key concerns (Baker 2018). To the extent that the bank data only cover a small fraction of the households' investments or the households in the bank sample are very different from other stock owners, the results would be difficult to interpret. The government data enable us to address both concerns.

To assess completeness, we compare the end-of-year values of the stock market portfolios we observe in Danske Bank to the same households' full stock market portfolios across all Danish financial institutions as stated on their tax return. The

[5] We use mapping files from Statistics Denmark to group granular education and industry codes into higher-level categories (Statistics Denmark 2017a, b).

portfolio data from Danske Bank is generally almost complete with only the largest investors occasionally holding stocks through other banks (see online Appendix Figure A2).

To assess representativeness, Table 1 compares the population of stock owners in Denmark (column 1) to the subsample who are active Danske Bank customers (column 2) and the subsamples who further survive the sample restrictions (columns 3 and 4) in terms of sociodemographics, balance sheets, and income. The households in the estimation sample are generally very similar to the population of stock owners although they are a few years older and have somewhat more net wealth and slightly less labor income.

The average household in the estimation sample has total assets just below \$500,000, with financial assets accounting for around \$175,000 and just below half of that, around \$80,000, invested in the stock market. Monthly gains from the stock market investments are around \$600, around one-third dividends and two-thirds price gains. The mean portfolio return in the estimation sample follows the main Danish stock index closely (see online Appendix Figure A3) although it is less volatile due to diversification with foreign investments. There is significant variation in portfolio returns across stockholders within each month: the difference between the fifth and ninety-fifth percentile of portfolio returns in a month is generally larger than 10 percentage points and as large as 40 percentage points in periods with market turmoil.

II. Empirical Framework

A. *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 naïve model:[6]

$$C_{i,t+1} - C_{i,t} = \alpha_t + \beta G_{i,t+1} + \epsilon_{i,t+1}, \tag{2}$$

where β expresses the marginal propensity to consume out of wealth gains. The naïve 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 endogeneous 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

[6] We follow most papers in relating ΔC to ΔW (Dynan and Maki 2001; Paiella and Pistaferri 2017; Aladangady 2017). Some papers relate $\Delta \log(C)$ to $\Delta \log(W)$ and estimate elasticities rather than MPCs (Mian, Rao, and Sufi 2013), while others relate ΔC to ΔG (Di Maggio, Kermani, and Majlesi 2020).

Table 1—Summary Statistics

	All stock owners (1)	Stock owners who are active DB customers (2)	Stock owners who are active DB customers with portfolio at the bank and not only DB stocks (3)	Stock owners who are active DB customers with portfolio at the bank, not only DB stocks, and no stocks in employer company (4)
Demographics				
Age of oldest member (years)	56.7	57.8	58.3	59.4
Highest education				
High school	0.56	0.51	0.50	0.51
College	0.29	0.29	0.29	0.29
Graduate school	0.14	0.18	0.19	0.18
Two-adult household	0.61	0.55	0.54	0.52
DB exclusive customer	0.14	0.52	0.56	0.57
End-of-year balance sheet (tax data)				
Total assets ($)	417,302	438,395	473,357	469,567
Financial assets ($)	120,645	145,231	171,286	174,118
Risky share	0.36	0.42	0.48	0.48
Total debt ($)	181,274	169,472	166,878	155,639
Income (tax data) and spending (DB data)				
Total income ($ per month)	7,228	7,555	7,563	7,042
Labor income ($ per month)	5,229	5,395	5,312	4,620
Spending ($ per month)		4,186	4,252	4,097
DB stock market portfolio (DB data)				
Portfolio size ($)			78,274	80,471
Portfolio risk			0.049	0.046
Expected monthly return			0.009	0.008
Number of securities			4.2	4.3
Actual gains ($ per month)			600	583
Passive gains ($ per month)			608	591
Dividends ($ per month)			173	180
Portfolio links to labor market				
Share of portfolio in own industry			0.09	0.05
Any equity in own industry			0.16	0.10
Any unlisted assets in portfolio			0.14	0.11
Number of households	1,648,654	451,463	366,129	354,282
Number of household-month observations	90,201,405	20,862,871	14,799,523	13,451,485

Notes: The table shows summary statistics for various samples of stock-owning households. Stock ownership is defined as having a portfolio of at least DKK 100 (about US$15). A household is an active Danske Bank (DB) customer in a given month if its adult members have at least one spending transaction each, and at least five jointly. Column (4) corresponds to our baseline estimation sample. End-of-year balance sheet measures are based on data from tax records: "Total assets" include financial assets and housing. "Financial assets" include securities and bank deposits. "Risky share" is the ratio of risky assets to total financial assets. "DB stock market portfolio" measures are based on data from DB: "Portfolio size" is the dollar value of the portfolio. "Portfolio risk" is the standard deviation of the current portfolio's monthly returns over the past 12 months. "Expected monthly return" is the mean of the current portfolio's monthly returns over the past 12 months. "Passive gains" are the dollar gains earned on the portfolio held at the beginning of the month. All Danish krone amounts are converted to US dollars using the average exchange rate from 2009 to 2016 of DKK 5.85 equaling USD 1. Table A1 in the online Appendix provides further distributional statistics for each variable.

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.[7]

We address this type of endogeneity by controlling for the portfolio characteristics with a rich set of nonparametric 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 12 months; and 20 indicators of *expected return* based on a ranking of households by the mean return of this portfolio over the past 12 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 that have different stock market gains in a given month despite entering the month with portfolios that have similar characteristics. The variation in gains exists because portfolios with similar characteristics are generally composed of different individual stocks with 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 that enter a period with 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. This introduces an endogeneity problem as buying and selling may be shaped by a range of shocks that also influence spending decisions. It may even take the form of reverse causality if households sell stocks to raise cash for consumption.

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 (see below).

We note that, theoretically, spending responses to stock market gains may depend on whether gains are created by shocks to expected firm profits or by shocks to the discount rate. Arguably, the time fixed effects in our empirical model absorb shocks to the discount rate, suggesting that the residual variation that identifies our estimates is created by idiosyncratic shocks to expected profits.

B. *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 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

[7] While standard portfolio theory predicts a unique correspondence between risk and expected return, risk may vary conditional on expected return when some households are off the efficient frontier.

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 hold portfolios that perform better conditional on past returns, in which case the assignment of the instrument would not be conditionally random.

To assess this threat to identification, we conduct a type of test adopted from the literature on lotteries (e.g., Cesarini et al. 2016; Cesarini et al. 2017).[8] 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 four investor characteristics. As shown in Figure 1, there is a strong *unconditional* correlation between investor characteristics and passive gains in all dimensions: investors who are older, hold more deposits, have higher monthly spending, and invest a larger share of their financial wealth in the stock market tend to have larger stock market gains. However, consistent with the identifying assumption, these correlations vanish almost completely when we introduce portfolio controls.[9]

C. *Dynamics*

We capture the fully dynamic spending response to stock market gains by considering 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.[10] We thus estimate the following model separately for different values of h:

$$C_{i,t+h} - C_{i,t} = \alpha_t + \beta G_{i,t+1} + \Theta_{i,t} + \gamma X_{i,t} + \mu_{i,t+h}, \tag{3}$$

where the actual stock market gains $G_{i,t+1}$ are instrumented with the passive gains, $\Theta_{i,t}$ captures portfolio characteristics as before, and $X_{i,t}$ is a set of household characteristics included to reduce the residual variation: age, education, and the number of children in the household.

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

[8] Section III considers another diagnostic: spending trends in the preperiod.

[9] In most but not all cases, the differences between the means illustrated in Figure 1 are not just economically immaterial but also statistically insignificant despite the large sample size.

[10] Jordà (2005) pioneered this local projections methodology.

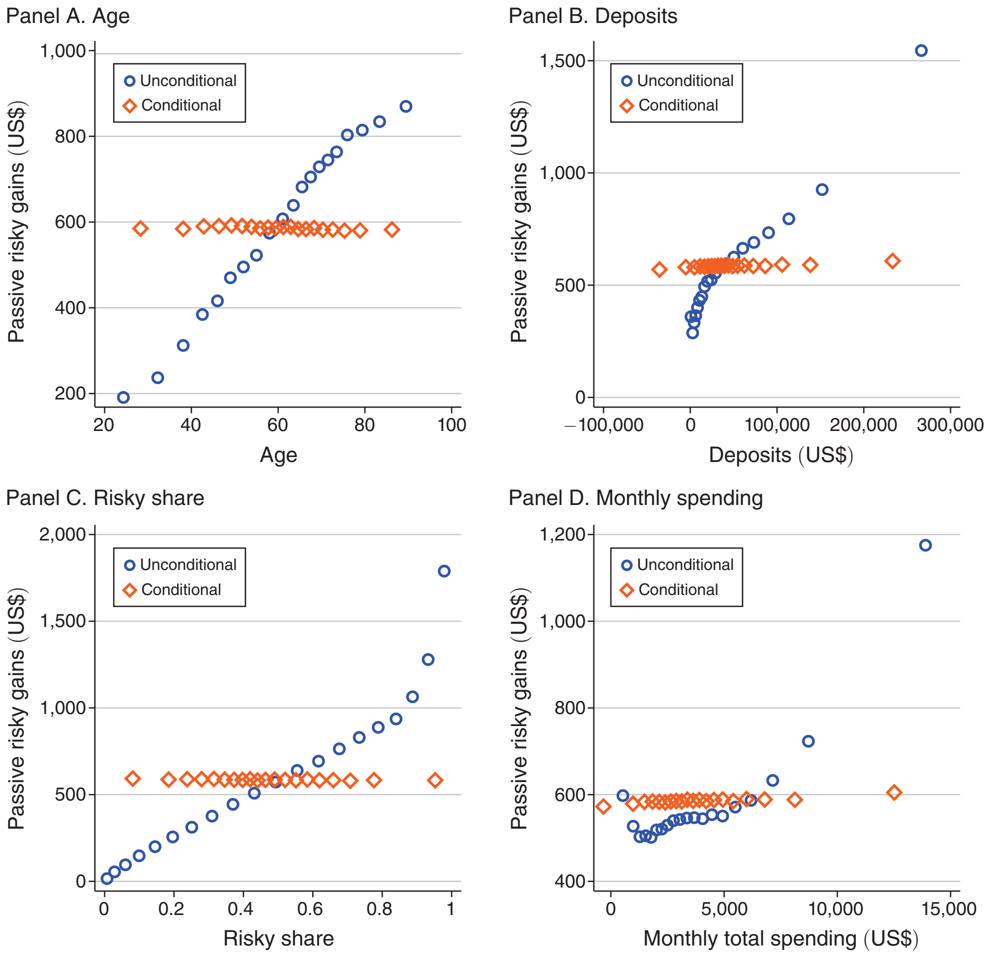

Figure 1. Lottery Tests

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

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 preperiod: 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 preperiod would be indicative of endogeneity.

The dynamic model requires us to consider the possibility that gains are serially correlated (conditional on controls) and that our estimates for period $t + h$ are identified not just by differences in gains in $t + 1$ but also by differences in gains in $t + 2$ and onward. On the one hand, gains are naturally correlated over time due to compounding even when the underlying returns are uncorrelated. The effect of compounded gains is part of the causal effect of gains in $t + 1$, which is what we want to estimate. On the other hand, serially correlated gains may also reflect a correlation in the underlying returns—for example, momentum (Fama and French 2012). In that case, our estimates for $t + h$ partly reflect differences in gains after $t + 1$ that are not caused by differences in gains in $t + 1$.

We investigate this issue by regressing actual gains in $t + h$ on passive gains in $t + 1$ conditional on the controls (see online Appendix Figure A4). For $h = 1$, this is precisely the first stage of (3), and the highly significant coefficient of 0.98 confirms that the instrument is relevant. For $h > 1$, the estimates are generally close to zero but more often positive than negative, suggesting that gains are indeed slightly correlated over time due to either compounding or serially correlated returns. In a robustness test, we eliminate the potentially confounding effect of serially correlated returns by controlling for future returns—that is, the average return between $t + 2$ and $t + h$ of the portfolio held at the beginning of $t + 1$.

III. Results

A. *Main Results*

In Figure 2, we illustrate the estimated dynamic spending responses to a \$1 gain. 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 preperiod. 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 in every month throughout the three-year estimation window.

In Table 2, we further provide estimates of the cumulative marginal propensities to consume over different time horizons.[11] After six months, investors have increased cumulative spending by 1.9 cents for each \$1 gain, rising to 4.4 cents after one year, 10.5 cents after two years, and 16.4 cents after three years. Hence, today's stock market gains feed into future household consumption at a rate of 4.4 percent in the first year and around 6 percent in the second and third years.

The finding that households increase consumption by a small amount in the present and all future periods is broadly consistent with the permanent income

[11] The cumulative estimates are not identical to the sum of the estimates in Figure 2. See table notes for details.

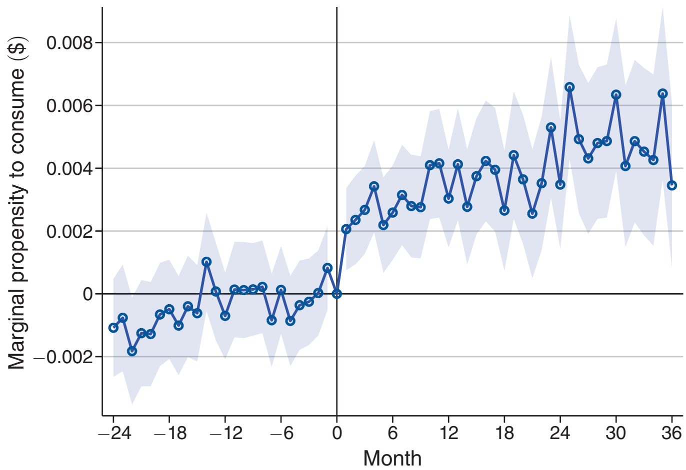

Figure 2. Marginal Propensity to Consume

Notes: The figure illustrates our dynamic estimates of the marginal propensity to consume: the effect of a $1 stock market gain in month $t+1$ on spending in month $t+h$. 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.

hypothesis. Given the estimated marginal propensity to consume and considering that stock market gains compound, the level shift in spending may in fact be permanent. This contrasts with the buffer stock behavior found in the context of lottery prizes (Fagereng et al. 2020) and unexpected inheritances (Druedahl and Martinello 2022), where households exhibit a much larger marginal propensity to consume and revert to the initial level of wealth. The discrepancy could reflect that the average owner of stock market wealth is highly liquid and therefore not constrained in the way that generates buffer stock behavior. We further explore the role of liquidity in the heterogeneity analysis below. It is also noteworthy that the adjustment to a higher level of monthly spending is somewhat sluggish, as evidenced by the upward drift in the point estimates in Figure 2, suggesting that the spending response is initially muted by a friction. We explore the role of a specific type of friction, inattention, in the heterogeneity analysis below.

B. *Robustness*

We probe the robustness of the results by modifying the estimating equation and the sample. First, we drop the requirement that households hold no investments in the firms where they work. Second, we restrict the sample to exclusive customers at Danske Bank to address concerns about secondary banks. Third, we exclude investments in firms in the industry of employment to eliminate confounding effects through changes in wage expectations. Fourth, we estimate the model without the household-level controls age, household size, and education. Fifth, we exclude investors holding unlisted shares to eliminate any confounding effect of unobserved gains from such investments. Finally, we eliminate the potentially confounding effect of

Table 2—Cumulative Marginal Propensity to Consume

	3 months	6 months	12 months	24 months	36 months
Main results					
Baseline	0.0084	0.0191	0.0442	0.1048	0.1641
	(0.0019)	(0.004)	(0.0083)	(0.0207)	(0.0382)
Robustness results					
Full sample	0.0094	0.0214	0.0482	0.1113	0.1542
	(0.0018)	(0.0037)	(0.0076)	(0.0188)	(0.0348)
Exclusive customers	0.0106	0.0288	0.0664	0.1232	0.1669
	(0.0029)	(0.0059)	(0.0125)	(0.0316)	(0.059)
Excluding own industry	0.0087	0.0200	0.0472	0.1227	0.1771
	(0.0021)	(0.0043)	(0.0091)	(0.0229)	(0.0424)
No controls	0.0086	0.0192	0.0439	0.1073	0.1687
	(0.0019)	(0.004)	(0.0083)	(0.0207)	(0.0383)
Excluding unlisted	0.0075	0.0185	0.0498	0.1206	0.1826
	(0.0021)	(0.0043)	(0.0089)	(0.0222)	(0.0410)
Future returns	0.0074	0.0205	0.0461	0.112	0.1717
	(0.0021)	(0.0042)	(0.0088)	(0.0222)	(0.0410)

Notes: 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). The estimate for horizon T is obtained by using as dependent variable $(C_{i,t+1} - C_{i,t}) + (C_{i,t+2} - C_{i,t}) + \ldots + (C_{i,t+T} - C_{i,t})$. "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 customers" means that households with accounts in banks other than Danske Bank are excluded from the estimation sample. "Excluding 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. "Excluding unlisted" means that we drop households with unlisted shares. Standard errors are clustered at the household level. We note that the cumulative MPC estimates are not identical to the sum of the monthly MPC estimates from the dynamic model as the estimation 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.

serially correlated portfolio returns by including three-way interactions between portfolio size at the beginning of month $t + 1$ (100 bins), the average return of this portfolio from month $t + 2$ to $t + h$ (20 bins), and calendar months. The resulting estimates of cumulative spending responses, reported in Table 2, are comparable to the baseline across all these alternative samples and specifications. The dynamics also consistently mirror the baseline with roughly parallel trends in the preperiod and a pronounced shift in the month of the wealth shock (see online Appendix Figure A5).

C. *Heterogeneity*

We investigate the heterogeneity in spending responses in four dimensions: liquidity, age, absolute return, and financial attention. In each dimension and within each month, we split the sample into three equally sized groups. There is significant variation in all dimensions: the average ratio of liquid assets to monthly income ranges from around 3 in the bottom tercile to more than 75 in the top tercile; financial attention ranges from 1 monthly log-in to around 13; age ranges from 40 years to almost 80; and the absolute monthly portfolio return ranges from 0.6 percent to 7.7 percent (see online Appendix Table A2). Figure 3 illustrates the cumulative spending responses for

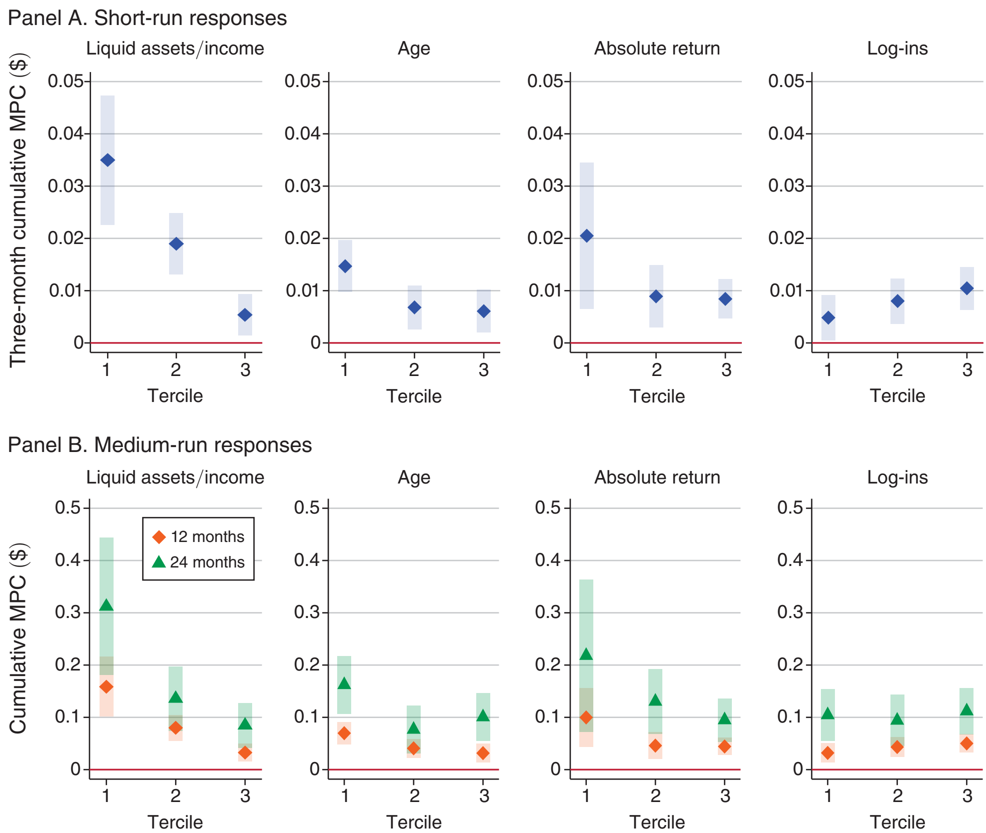

Figure 3. Heterogeneity in the Marginal Propensity to Consume

Notes: The figure illustrates the heterogeneity in estimated marginal propensities to consume out of stock market gains cumulated over a three-month horizon (panel A) and one-year and two-year horizons (panel B). There are four dimensions of heterogeneity: *liquidity* is beginning-of-month liquid assets relative to average monthly income; *age* is the beginning-of-month age of the oldest household member; *absolute return* is the absolute size of the portfolio return in the month; and *log-ins* is the number of days with log-ins to Danske Bank accounts in an average month. 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.

each group over the short run (three months) and medium run (one and two years), while the analogous dynamic estimates are illustrated in online Appendix Figure A6.[12]

There is a clear liquidity gradient in the MPC estimates both in the short and the medium run. Low-liquidity households spend 16 cents of a $1 gain over one year and 30 cents over two years, about four times more than high-liquidity households. 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). While it implies that a large group of stock owners consume their gains quickly, most households in this group hold small portfolios and the implications for aggregate consumption are modest.

[12] The two panels have different scales on the y-axis.

There is also a clear attention gradient, at least in the short run. While even the least attentive households exhibit statistically significant responses to stock market wealth shocks over three months, their cumulative MPC is only half as large as for the most attentive households. Over longer horizons, the attention gradient becomes less pronounced, and the two-year cumulative MPC is roughly equalized across groups with high and low financial attention. The results suggest that inattention plays a role in attenuating household responses to stock market cycles over horizons as long as one year.

Contrary to the prediction of standard life cycle models where older households consume more of their gains than younger households because their remaining lifespan is shorter and uncertainty about lifetime income largely has been resolved, we find slightly lower MPC estimates for older households. Several empirical papers have found a qualitatively similar relation between MPCs and age outside the domain of stock markets (e.g., Fagereng, Holm, and Natvik 2021). While the age gradient partly reflects differences in wealth—that is, older households tend to hold larger portfolios—the young continue to have the highest MPC when comparisons are made within wealth groups (see online Appendix Figure A7).

Finally, we find that households consume more of their gains when returns are closer to zero. A possible interpretation is that the welfare cost of failing to smooth is increasing in the absolute size of the gain. 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 (Kueng 2018).

IV. Conclusion

In this paper, we study how shocks to stock market wealth affect consumer spending over the short and the medium term. 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.

Our main finding is 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 rate of 4.4 percent in the first year and around 6 percent in the next years. The findings are broadly 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, although there is some evidence that the spending adjustment is initially dampened by financial inattention.

REFERENCES

Aladangady, Aditya. 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415–46.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman. 2019. "Tax Evasion and Inequality." *American Economic Review* 109 (6): 2073–103.

Andersen, Asger Lau, Amalie Sofie Jensen, Niels Johannesen, Claus Thustrup Kreiner, Søren Leth-Petersen, and Adam Sheridan. 2023. "How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Data Sets." *American Economic Journal: Applied Economics* 15 (4): 1–29.

Andersen, Asger Lau, Niels Johannesen, and Adam Sheridan. 2024. *Data and Code for: "Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi Lotteries on the Stock Market."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E196641V1.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model." NBER Working Paper 26647.

Baker, Scott R. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy* 126 (4): 1504–57.

Bräuer, Konstantin, Andreas Hackethal, and Tobin Hanspal. 2022. "Consuming Dividends." *Review of Financial Studies* 35 (10): 4802–57.

Bureau van Dijk. 2018. *ISIN-Firm Linkages*. Brussels, Belgium: Bureau van Dijk Electronic Publishing. Accessed May 7, 2018.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2007. "Down or Put: Assessing the Welfare Costs of Household Investment Mistakes." *Journal of Political Economy* 115 (5): 707–47.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2009. "Fight or Flight? Portfolio Rebalancing by Individual Investors." *Quarterly Journal of Economics* 124 (1): 301–48.

Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–51.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1): 1–55.

Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek. 2011. "How Large Are Housing and Financial Wealth Effects? A New Approach." *Journal of Money, Credit and Banking* 43 (1): 55–79.

Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling. 2017. "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries." *American Economic Review* 107 (12): 3917–46.

Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Wallace. 2016. "Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players." *Quarterly Journal of Economics* 131 (2): 687–738.

Chetty, Raj, and Adam Szeidl. 2016. "Consumption Commitments and Habit Formation." *Econometrica* 84 (2): 855–90.

Chodorow-Reich, Gabriel, Plamen T. Nenov, and Alp Simsek. 2021. "Stock Market Wealth and the Real Economy: A Local Labor Market Approach." *American Economic Review* 111 (5): 1613–57.

Cieslak, Anna, and Annette Vissing-Jorgensen. 2021. "The Economics of the Fed Put." *Review of Financial Studies* 34 (9): 4045–89.

Danske Bank. 2018. *Data on Customer Transactions, Account Balances and Financial Asset Portfolios, 2009–16*. Copenhagen, Denmark: Danske Bank. Accessed March 19, 2018.

Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi. 2020. "Stock Market Returns and Consumption." *Journal of Finance* 75 (6): 3175–219.

Druedahl, Jeppe, and Alessandro Martinello. 2022. "Long-Run Saving Dynamics: Evidence from Unexpected Inheritances." *Review of Economics and Statistics* 104 (5): 1079–95.

Dynan, Karen E., and Dean M. Maki. 2001. "Does Stock Market Wealth Matter for Consumption?" FEDS Working Paper 2001-23.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso. 2017. "Asset Market Participation and Portfolio Choice over the Life-Cycle." *Journal of Finance* 72 (2): 705–50.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1): 115–70.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13 (4): 1–54.

Fama, Eugene F., and Kenneth R. French. 2012. "Size, Value, and Momentum in International Stock Returns." *Journal of Financial Economics* 105 (3): 457–72.

Friedman, Milton. 1957. "The Permanent Income Hypothesis." In *A Theory of the Consumption Function*, edited by Milton Friedman, 20–37. Princeton, NJ: Princeton University Press.

Gabaix, Xavier. 2019. "Behavioral Inattention." In *Handbook of Behavioral Economics: Applications and Foundations*, Vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 261–343. Amsterdam: North-Holland.

Ganong, Peter, Damon Jones, Pascal J. Noel, Fiona E. Greig, Diana Farrell, and Chris Wheat. 2020. "Wealth, Race, and Consumption Smoothing of Typical Income Shocks." NBER Working Paper 27552.

Gargano, Antonio, and Alberto G. Rossi. 2018. "Does It Pay to Pay Attention?" *Review of Financial Studies* 31 (12): 4595–649.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. "Harnessing Naturally Occurring Data to Measure the Response of Spending to Income." *Science* 345 (6193): 212–15.

Harvey, Campbell R., and Akhtar Siddique. 2000. "Conditional Skewness in Asset Pricing Tests." *Journal of Finance* 55 (3): 1263–95.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–610.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy according to HANK." *American Economic Review* 108 (3): 697–743.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–239.

Kaplan, Greg, and Giovanni L. Violante. 2018. "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 32 (3): 167–94.

Karlsson, Niklas, George Loewenstein, and Duane Seppi. 2009. "The Ostrich Effect: Selective Attention to Information." *Journal of Risk and Uncertainty* 38 (2): 95–115.

Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics* 133 (4): 1693–751.

Lettau, Martin, and Sydney C. Ludvigson. 2004. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." *American Economic Review* 94 (1): 276–99.

Markowitz, Harry. 1952. "The Utility of Wealth." *Journal of Political Economy* 60 (2): 151–58.

Meyer, Steffen, and Michaela Pagel. 2019. "Fully Closed: Individual Responses to Realized Gains." *Journal of Finance* 77 (3): 1529–85.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–726.

Olafsson, Arna, and Michaela Pagel. 2017. "The Ostrich in Us: Selective Attention to Financial Accounts, Income, Spending, and Liquidity." NBER Working Paper 23945.

Paiella, Monica, and Luigi Pistaferri. 2017. "Decomposing the Wealth Effect on Consumption." *Review of Economics and Statistics* 99 (4): 710–21.

Poterba, James M. 2000. "Stock Market Wealth and Consumption." *Journal of Economic Perspectives* 14 (2): 99–118.

Reis, Ricardo. 2006. "Inattentive Consumers." *Journal of Monetary Economics* 53 (8): 1761–800.

Sicherman, Nachum, George Loewenstein, Duane J. Seppi, and Stephen P. Utkus. 2016. "Financial Attention." *Review of Financial Studies* 29 (4): 863–97.

Statistics Denmark. 2017a. *Education Codes*. Copenhagen, Denmark: Statistics Denmark, Research Service. Accessed November 30, 2017.

Statistics Denmark. 2017b. *Industry Codes*. Copenhagen, Denmark: Statistics Denmark, Research Service. Accessed November 30, 2017.

Statistics Denmark. 2018a. *Befolkningen*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/da/Statistik/emner/borgere/befolkning/befolkningstal (accessed July 31, 2018).

Statistics Denmark. 2018b. *Beskæftigelse for lønmodtagere*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/da/Statistik/dokumentation/statistikdokumentation/beskaeftigelse-for-loenmodtagere--md--/indhold (accessed July 31, 2018).

Statistics Denmark. 2018c. *IDAP - IDA persondata*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/IDAP%20-%20IDA%20persondata.html (accessed July 31, 2018).

Statistics Denmark. 2018d. *ILME - Ikke lønmodtagerdata fra E-Indkomst - akkumuleres til års-register over de 4 kvartaler*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/ILME%20-%20Ikke%20l%C3%B8nmodtagerdata%20fra%20E-Indkomst%20-%20akkumuleres%20til%20%C3%A5rs-register%20over%20de%204%20kvartaler.html (accessed July 31, 2018).

Statistics Denmark. 2018e. *Indkomst*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/IND%20-%20Indkomst.html (accessed July 31, 2018).

Statistics Denmark. 2018f. *Uddannelser*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/UDDA%20-%20Uddannelser%20(BUE).html (accessed July 31, 2018).
Statistics Denmark. 2021a. *IRTEPERS - Indlånsrenter - Kontospecifikke data for personer*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/IRTEPERS%20-%20Indl%C3%A5nsrenter%20-%20Kontospecifikke%20data%20for%20personer.html (accessed November 23, 2021).
Statistics Denmark. 2021b. *URTEPERS - Udlånsrenter - Kontospecifikke data for personer*. Copenhagen, Denmark: Statistics Denmark, Research Service. https://www.dst.dk/extranet/ForskningVariabellister/URTEPERS%20-%20Udl%C3%A5nsrenter%20-%20Kontospecifikke%20data%20for%20personer.html (accessed November 23, 2021).