ARTICLE IN PRESS

Journal of Monetary Economics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Wealth shocks and portfolio choice[☆]

Dimitris Christelis[a], Dimitris Georgarakos[b], Tullio Jappelli[c,*], Geoff Kenny[b]

[a] *University of Glasgow, CSEF, CFS, and Netspar, United Kingdom*
[b] *European Central Bank, Germany*
[c] *University of Naples Federico II, CSEF and CEPR, Italy*

ARTICLE INFO

JEL Classification:
D14
G11
G51

Keywords:
Household finance
Stock market participation
Risk aversion
Consumer Expectations Survey

ABSTRACT

We use new euro area representative data from the Consumer Expectations Survey (CES) to elicit household-specific propensities to invest and consume out of positive wealth shocks. Using a randomized assignment of hypothetical lottery gains ranging from €5,000 to €50,000 and a realistic menu of consumption, saving and asset choices, we estimate the causal effect of wealth shocks on risky asset ownership and conditional asset shares. Wealth shocks have a positive effect on stockholding (between 8.4 and 12.8 percentage points increase in participation for the largest wealth shock). The majority of households do not participate in the stock market, even after a large increase in wealth. The conditional asset share invested in risky assets is constant for wealth shocks up to €20,000, and edges up slightly (by at most 2 %) for larger prizes. Our evidence is consistent with constant relative risk aversion for the majority of risky asset investors, while we also find important heterogeneity in the level of risk aversion across individuals.

1. Introduction

The relation between household resources and portfolio decisions has been a subject of intense empirical investigation in the last two decades. This relation speaks to the importance of transaction and information costs for accessing financial markets, as well as to the nature of attitudes towards risk with resulting implications for the transmission of financial shocks both in the long-run and over the business cycle.

However, deriving causal estimates of the effect of wealth on stock market participation and the share of resources invested in risky assets is challenging when using observational data, and even longitudinal ones.[1] Causality could run from participation and amounts

☆ The opinions expressed in the paper are those of the authors and do not necessarily reflect the views of the European Central Bank or the euro system. We thank Boragan Aruoba (the Editor), Luminita Stevens (the Associate Editor) and one anonymous referee, Luigi Guiso, Morten Ravn and Jirka Slacalek for helpful comments. Dimitris Christelis and Tullio Jappelli acknowledge financial support from the European Savings Institute (OEE). Jappelli also acknowledges partial financial support from the EU in the framework of the GRINS project (PE00000018 – CUP E63C22002140007). We thank participants at the joint ECB-NY FED conference on Expectations Surveys, the Royal Economic Society Conference, Bank of Italy Conference on "The use of surveys for monetary and economic policy" (Rome, 26–27 April 2023) and at seminars at the EEA-ESEM 2023 Congress, the OEE and the ECB.

* Corresponding author at: Department of Economics and Statistics, University of Naples Federico II, Via Cinzia, Monte S. Angelo, Napoli 80126, Italy.

E-mail address: tullio.jappelli@unina.it (T. Jappelli).

[1] See for instance Chiappori and Paiella (2011), Brunnermeier and Nagel (2008), Calvet et al. (2009) and the survey by Guiso and Sodini (2013).

https://doi.org/10.1016/j.jmoneco.2024.103632
Received 21 June 2023; Received in revised form 15 July 2024; Accepted 15 July 2024

invested in risky assets to resources, as households that participate in financial markets may hold more diversified portfolios and be successful in securing high returns, thus ending up with higher wealth.[2] In addition, there could be a correlation between wealth and household (time-varying) unobserved characteristics that affect the decision to invest in risky financial assets. For instance, wealth and risky asset investment could be both correlated with unobservable investment preferences, such as patience and the length of the investment horizon. Therefore, any observed positive correlation between risky asset investment and wealth may be driven by associations with unobserved factors or might reflect the fact that people who have chosen to participate in the stock market independently of their wealth also end up with higher wealth.

To address these issues, some studies have measured the effect of exogenous wealth shocks on stockholding through particular events such as inheritances in Denmark (Andersen and Nielsen, 2011) and lottery prizes in Sweden (Briggs et al., 2021). However, there could be questions about the external validity of these results, as a sample consisting of lottery winners and those who receive inheritances is unlikely to be representative of the general population. Moreover, both these studies focus on the extensive margin of stock investing.

To overcome the challenges affecting observational studies, we implement a novel experiment in a large-scale household survey that is representative of the euro area population. We ask respondents how they would allocate randomly assigned lottery gains of different size across spending, debt repayment, and investment in various financial assets. This set-up allows us to estimate the causal effect of wealth shocks on both stock market participation and the share of financial assets allocated into stocks and mutual funds (i.e., the intensive margin).

Our experimental approach extends a fast-growing literature that has used questions in household surveys to elicit information about the consumption responses to scenarios involving income or wealth shocks.[3] A key advantage of this approach is that it provides household-specific estimates of the marginal propensity to consume out of transitory income changes and wealth shocks. By suitable experimental designs, these estimates apply to the population at large, allow comparisons under different shock scenarios, and are not confounded by the unobserved characteristics of the selected sample that is subject to a realized shock.

We use data for about 10,000 households from the six largest euro area economies (Germany, France, Italy, Spain, the Netherlands and Belgium) from the Consumer Expectations Survey (CES), an ongoing panel administered by the ECB. This survey, which is described in Georgarakos and Kenny (2022) and ECB (2021), is representative of the national populations and collects via the Internet fully harmonized information on demographics, income, consumption, portfolios, and several expectations variables. In June 2021, we implemented a special set of questions on the way that households would spend, save, and invest windfall gains from lotteries of five randomly assigned amounts that vary from €5000 to €50,000.

Given that we randomize the lottery prizes, we estimate the causal effect of the lottery-induced changes in wealth on stock market participation and the share of resources invested in risky financial assets. The design of our experiment allows us to derive causal and population-representative estimates of the effect of wealth on these two margins, and to examine responses of various population groups. In the experiment, the wealth shock is exogenous, and not the result of the amount saved in previous periods, or of the portfolio allocation chosen by investors in the past. Therefore, we address the issue that portfolio choice may spuriously respond to endogenous changes in wealth, induced for instance by the fact that some investors have been successful in their asset allocation, avoiding market crashes or exploiting market booms.

Our results can be summarized as follows. On participation, we confirm previous literature reporting a positive gradient between wealth shocks and stock ownership, see Guiso and Sodini (2013) for a review. Winning the highest (relative to the lowest) prize, implies an increase in participation between 8.4 and 12.8 percentage points, depending on the definition of stockholding. Even for relatively large wealth shocks (€50,000) the majority of investors chooses not to invest in the stock market. The risky asset share is rather insensitive to small or moderate wealth shocks (up to €20,000) and increases by at most 2 percentage points for larger ones (€30,000 or €50,000). The result applies to investors with different levels of wealth, financial sophistication, trust, income, and stock market expectations. As we argue, our evidence is consistent with constant relative risk aversion (CRRA) for the majority of risky asset investors.

The paper is organized as follows. Section 2 reviews briefly the relevant household finance literature. Section 3 describes the survey and the experimental framework. Section 4 provides a first description of the data. Section 5 presents econometric estimates of the participation decisions in stocks and other risky assets, and of the asset share invested in stocks. Section 6 compares actual and hypothetical data in the CES. Section 7 discusses the implications for the macro literature and Section 8 concludes.

2. Stockholding, portfolio choice and wealth

A clear finding across many studies and countries is that participation in stockholding is stronger among wealthier households, and that there are often significant differences in participation rates between households across the wealth distribution. The positive relation between stockholding and wealth is largely interpreted as evidence that fixed entry costs, information costs, minimum investment requirements, and participation fees limit severely stock market participation (Guiso and Sodini, 2013).

[2] In a different context, Carroll et al. (2011) explains the importance of exogenous wealth shocks pointing out that identifying the "pure" housing wealth effect is hard, because "one would want data on spending by individual households before and after some truly exogenous change in their house values, caused for example by the unexpected discovery of neighborhood sources of pollution."

[3] See contributions by Shapiro and Slemrod (2003), Jappelli and Pistaferri (2014), Christelis et al. (2019). Stantcheva (2023) discusses the benefits of asking scenario questions in household surveys.

Studies also find that many households do not participate in the stock market even at high levels of wealth and that, at any given level of wealth, households with higher financial sophistication are more likely to participate (van Rooij et al., 2011). More sophisticated households are able to process information more efficiently and more cheaply, and therefore face lower entry and participation costs than less sophisticated households, for any given level of wealth. Other studies suggest that non-participation, especially among the wealthy, is due to lack of trust in financial markets and pessimistic beliefs about stock market returns Guisoet al. (2008), and/or psychological barriers, such as status quo bias, loss aversion, and present bias.

Another strand of literature focuses on portfolio allocations in various asset categories. Portfolio shares (including non-participants) do not distinguish the effects of relevant variables on the participation decision from those on the portfolio share, conditional that the risky asset is held. The literature has therefore modelled portfolio shares, considering that many households do not invest in risky assets. Econometrically this amounts to estimating models that allow for sample selection.

Estimating the relation between wealth and the asset share is useful to test some of the implications of portfolio models. A classic prediction of these models is that in the absence of human capital and when preferences are characterized by CRRA, the portfolio share of risky assets, conditional on holding them, should be independent of the level of wealth.[4] Several recent papers show that an estimated constant wealth share of risky assets does not necessarily imply CRRA, unless one controls fully for a number of (sometimes unobserved) confounding factors, such as income risk, habits, human wealth, and committed expenses. For instance, Wachter and Yogo (2010) suggest that the riskiness of human capital affects optimal portfolio choice, so that the proportion of human capital to financial wealth is a confounding factor in the relation between financial wealth and the share of risky assets. Meeuwis (2022) shows that positive income shocks and increases in financial wealth affect, in opposite directions, the risky asset share. Calvet and Sodini (2014) point out that income risk, habits, committed expenses and other variables (such as real estate wealth, family composition, and borrowing constraints) can obscure the relation between the risky asset share and wealth.

One approach to address the challenge of the possible confounding factors and obtain causal estimates of the effect of wealth on risky assets shares is to use panel data, as in Chiappori and Paiella (2011). These authors find that the (log) change in the risky asset share is rather insensitive to (log) change in wealth, as suggested by utility functions with CRRA. Using panel data is only a partial solution because this approach still requires distinguishing variations in asset shares that are due to active saving from variations due to changes in asset prices. Furthermore, past wealth is likely to be correlated with the error term, due to portfolio inertia, while time-varying unobservable variables that affect both asset shares and wealth can still lead to inconsistent estimates. Measurement error in wealth poses an additional problem, as it can lead to the usual attenuation bias of the estimated elasticity of the risky asset share with respect to wealth. In any case, even in panel data, one needs to identify wealth shocks that are exogenous to the portfolio decision.[5]

Calvet and Sodini (2014) have proposed an ingenious alternative to OLS and first difference estimates, controlling for unobserved fixed effects using Swedish data on the portfolios of twins. In contrast to Chiappori and Paiella, they find that the risky asset share increases with financial wealth, lending support to the notion that preferences are characterized by decreasing relative risk aversion (DRRA).[6]

A promising route to overcome the identification problems is to exploit random wealth shocks due to unexpected inheritances or lottery prizes. Andersen and Nielsen (2011) use Danish inheritances from sudden deaths and find that unexpected inheritances increase stock market participation by 21 percentage points within the following three years for those who receive €134,000, but that the majority of households do not enter the stock market even after receiving a large inheritance. Briggs et al. (2021) use administrative Swedish data exploiting randomly received lottery prizes and find that a $150,000 windfall gain increases the probability to enter the stock market by 12 percentage points among pre-lottery gain non-participants, but has no discernible effect on pre-lottery gain stockowners.

Our paper is closely connected to these two studies using administrative data. However, rather than using observational data, we rely on an experimental design that is robust to confounding effects and allows us to estimate the causal effect of exogenous wealth shocks on stockholding. Our experiment allows us to study the effect of windfall gains also on risky asset shares using population representative data for the six largest euro area countries, and to examine possible different effects of wealth shocks across various subgroups.

[4] When human capital is riskless and tradeable, it induces poorer households (those who have a lower ratio of financial wealth to labor income) to invest more in risky assets to take advantage of the equity premium, while richer households (with high financial wealth relative to labor income) should prefer to hold a higher share of riskless assets. This classical prediction of Merton (1971) in a complete market setting has been refined by many papers in life-cycle settings and incomplete markets (a prominent example is Cocco et al., 2005).

[5] In principle, one could also control directly for the propensity to take risk if this propensity were observable. This would provide information on risk aversion, but would of course not solve the causality problem.

[6] Brunnermeier and Nagel (2008) and Bilias et al. (2010), using PSID data, found risk taking to be insensitive to wealth changes, and that one of the major drivers of households' portfolio allocation seems to be portfolio inertia: households rebalance only very slowly following inflows and outflows or capital gains and losses. Calvet et al. (2009) find that in Sweden households rebalance toward a higher risky share as they become richer. Paravisini et al. (2017) using data from a lending platform find evidence for a positive relation between wealth and risk taking.

3. The consumer expectations survey and the experimental design

In our analysis, we use the ECB's Consumer Expectations Survey (CES), an online high frequency panel survey of euro area consumer expectations and behavior.[7] In this paper, we use mostly data from a special-purpose survey fielded in June 2021. The CES data during the period of our analysis covers the six largest euro area economies with a sample size of approximately 10,000 consumers: 2,000 survey participants from each of the four largest euro area countries (Germany, Italy, France, Spain), and 1,000 in each of the two next largest countries (Belgium, the Netherlands). The survey provides sample weights that we use to make descriptive statistics representative of the adult population in each country.

Respondents are invited to answer online questionnaires every month and then typically leave the panel between 18 and 24 months after joining. Each respondent completes a questionnaire upon entry into the panel, providing background information that changes very little month by month (e.g., education, family situation, measures of financial literacy). More time-sensitive information is collected in a series of monthly, quarterly, and ad hoc topical questionnaires.

In June 2021, we asked respondents in the CES to report how they would allocate a hypothetical lottery prize, distinguishing between spending, saving, and investing in financial assets, and repaying debt. The question randomly assigns five different lottery prizes (<*Amount*>: 5, 10, 20, 30 and 50 thousand euro) across respondents: *Imagine you win a lottery of* ⟨*Amount*⟩ *today. How would you use this unexpected windfall over the next 12 months? Please allocate the* ⟨*Amount*⟩ *over the following three categories.*

- Buy goods and services (including food, housing costs and rent, utility bills, clothing, and long-lasting goods such as home improvements, furniture and electronics, etc.).
- Save and invest in financial assets.
- Repay debts.

In a follow up question, respondents who indicated that they would save were asked: *You said that you would save or invest in financial assets € <Amount>. Please indicate in which of the following asset categories you would save/invest this amount.* Financial investment categories consist of (i) current accounts and saving accounts; (ii) stocks and shares; (iii) mutual funds; (iv) retirement and pension products (including life insurances); and (v) bonds. The Appendix reports all details of the survey questions.

A number of recent studies use survey data and quantitative responses to transitory income and wealth shock scenarios to elicit the marginal propensity to consume (e.g., Christelis et al., 2019 and Jappelli and Pistaferri, 2014, amongst others). We extend this literature by asking respondents to indicate their choices out of a realistic menu that includes not only consumption and saving, but also asset investment choices and portfolio allocation. Another novel feature of our approach is that it randomly assigns lottery prizes of different size to survey participants which enables the identification of the causal relationships of interest.

Three additional features of the survey questions are noteworthy. First, consumers are asked by how much they would consume or save "in the next 12 months". This allows us to rule out that differences in the amounts consumed, saved, and invested arise from differences in the timing of planned spending.[8] Second, the questionnaire was administered in June 2021. Even though the COVID-19 period cannot be regarded as a normal period, June 2,021 did not coincide with any of the major waves of the pandemic. Finally, since median financial wealth in the sample is €15,105, we experiment with a range of shocks whose amounts are realistic for median households. Indeed, the largest lottery (€50,000) amounts to roughly the financial wealth of respondents at the 75th percentile of the financial wealth distribution (€53,657 euro).

Our survey experiment is part of a large and growing literature on economic expectations and survey experiments. Bachmann et al. (2022) provide a comprehensive guide on the survey process, from the design of questions and experiments to the collection of data. Often the design involves asking respondents to make hypothetical decisions described in real life situations in which some factors vary experimentally. Stantcheva (2023) provides examples in which this approach has been used in different fields, such as education, labor, health, and macro-finance.

One concern of this approach is that individuals may have different incentives when responding to hypothetical questions compared to their actual behavior, for a variety of reasons. Some respondents may not perceive any personal benefit or consequence from their responses, leading to less reliable answers. Others may be inclined to provide answers that are socially desirable or acceptable or may be influenced by perceived societal norms (for instance, reporting high saving rates or propensity to repay debt) even if these responses do not reflect their actual behavior. Still others may consider financial behavior a private or sensitive issue, and may be reluctant to disclose it in a survey setting. It is also possible that hypothetical questions may elicit more candid responses because they don't require the same level of disclosure as reporting actual behavior.[9] Importantly, our randomization implies that these potential biases are randomly distributed across subsamples, so one can still make causal inference and obtain useful information from the experiment. Moreover, after presenting our main results, we compare our experimental approach with findings using observational data from different data sources and with observational portfolio data from the CES.

[7] Building on recent international experiences and advances in survey methodology and design, as reflected, for example, in the Federal Reserve Bank of New York's Survey of Consumer Expectations (Armantier et al., 2016), the CES was launched in pilot phase in January 2020.

[8] In principle, it would be useful to posit similar questions with other time horizons (e.g., how consumption, saving and financial investments would change in the following years), but this would considerably increase the complexity of the questionnaire.

[9] Since some respondents may have difficulty predicting their behavior in hypothetical situations, we conducted pilot testing to identify any issues with question wording or interpretation.

4. Descriptive analysis

We start by reporting sample statistics on the randomized questions and on the main variables used in the estimation. Means and standard deviations are computed using sample weights. To control for survey-related measurement error and noise, we exclude 538 respondents (approximately 5 % of the sample) who completed the survey (which for a typical respondent should take about 10 min) in less than 2.5 min, so that the resulting sample includes 9,668 observations.[10]

In Table 1, to compute the marginal propensity to consume (MPC), to save (MPS) and to repay debt (MPD), we standardize the reported amounts by the randomly assigned lottery prizes. On average, the MPC out of the wealth shocks is 34 %, the fraction of wealth used to repay debt is 18 %, and the remaining 48 % is saved. These estimates are broadly in line with similar propensities estimated for other countries with direct survey questions or actual windfall gains. Jappelli and Pistaferri (2010), Havranek and Sokolova (2020), Gelman (2021) and Crawley and Theloudis (2024) survey a large body of empirical evidence indicating that in the U.S. the average MPC on non-durable goods and services out of actual windfall gains of $500-$1,000 is between 15 % and 25 %. Jappelli and Pistaferri (2014) use data from the Italian Survey of Household Income and Wealth where consumers were asked to report the fraction of a positive income shock (a hypothetical unanticipated tax rebate) that they would save and find an average MPC of 48 %. Christelis et al. (2019), using similar questions, find that in the Netherlands the average MPC from nondurable consumption is 19 %. The average masks substantial heterogeneity, because many households have MPCs that are close to zero, while some have MPCs close to 1, with considerable variation in between.[11]

4.1. Allocation of the lottery prize between consumption and saving

Fig. 1 plots the MPC, the MPS and the MPD by the size of the lottery prize and associated 95 % confidence bands. In Table 2 we report regression results related to the same variables. The baseline specification includes four dummies for lottery gains (the omitted gain is €5,000) and country dummies. The extended specification controls also for demographic variables (gender, dummies for age groups, high school, college, employment status, and log cash-on-hand that is defined as monthly income plus financial wealth). We use Tobit models because by construction the MPC, MPS, and MPD are constrained to vary from 0 to 1; for instance, in the limit case that the entire prize is allocated to consumption, MPC equals one, and MPS and MPD are both equal to zero.

Standard intertemporal models with perfect markets predict that the MPC should not vary with the size of wealth shocks, while models with liquidity constraints predict a negative relation between MPC and shock size. If consumers receive a large and positive windfall, they are more likely to overcome the constraint, and therefore the associated MPC is lower than in the case of small gains. Consistent with this prediction, the regression results show that there is a negative relation between the MPC and lottery prizes: going from the smallest to the largest gain reduces the MPC by 4 percentage points. This evidence is in line with earlier studies (e.g., Christelis et al., 2019). The relation between MPC and shock size is still negative, but smaller in absolute value if we control for demographic variables and cash-on-hand in column (2) of Table 2. Results for MPD suggest no relation between shock size and the propensity to repay debt, while the relation with MPS is positive (between 7 and 9 percentage points higher for the largest shock relative to baseline). We take the evidence about the MPC as confirming that the experiment is highly informative, and responses consistent with previous evidence.

4.2. Asset allocation

The novelty of the survey experiment is that we also collect information on how respondents would allocate the amount saved (if positive) in various assets. We use a narrow and a broad definition of risky assets. The narrow definition includes only stocks. The broad definition includes also mutual funds and collective investments. Notice that this broad definition is likely above the combined amount of direct and indirect stockholding, as parts of mutual funds will also be invested outside the stock market. Even in this case, however, mutual funds are risker than transaction accounts.

Using this directly elicited information for asset allocation out of wealth shock scenarios we create dummy variables for direct and total stock market participation. These binary indicators do not distinguish between those who would increase stock investment (because they already have stocks), and those who would enter the market due to the lottery gain. For this, we use additional survey information on actual portfolios (collected prior to fielding the wealth shock question).

We also examine how risky asset shares change because of the lottery gains, and in addition whether respondents who already own stocks (and mutual funds) adjust their shares differently. For this purpose, we need to use actual portfolio data in combination with data from the hypothetical lottery and compute the change in the risky asset share after assigning the lottery gain. Defining period t-1 as the time immediately before the lottery question is asked, and period t as the time after the question is answered, the change in the asset share is:

$$\Delta\alpha_{it} = \frac{A_{it}}{W_{it}} - \frac{A_{it-1}}{W_{it-1}}, \tag{1}$$

[10] Results are almost identical if we keep these observations in the estimation sample.

[11] Gelman (2021) and Kaplan and Violante (2022) conduct a systematic investigation of intertemporal consumption models to study the relative importance of circumstances and characteristics in explaining MPC heterogeneity.

Table 1
Descriptive statistics.

	Mean	Std. Dev.	N
MPC (Marginal Propensity to Consume)	0.340	0.322	9677
MPS (Marginal Propensity to Save)	0.479	0.355	9677
MPD (Marginal Propensity to Repay Debt)	0.180	0.296	9677
Propensity to invest in stocks	0.257	0.437	9677
Propensity to invest in stocks or MF	0.380	0.485	9677
Participation in stocks before the lottery	0.232	0.422	9677
Participation in stocks or MF before the lottery	0.330	0.470	9677
Age 18–34	0.242	0.428	9677
Age 35–49	0.327	0.469	9677
Age 50–64	0.323	0.468	9677
Family size	2.718	1.271	9677
High school	0.319	0.466	9677
College	0.542	0.498	9677
Employed	0.587	0.492	9030
Self-employed	0.097	0.296	9030
High literacy	0.523	0.499	9565
High trust	0.530	0.499	9677
Liquidity	0.735	0.441	9677
Cash-on-hand	60.731	125.901	9149
Belgium	0.094	0.291	9677
Germany	0.207	0.405	9677
Spain	0.205	0.404	9677
France	0.194	0.395	9677
Italy	0.214	0.410	9677
Netherlands	0.087	0.282	9677
Share invested in stocks in the lottery (conditional)	0.402	0.296	2491
Share invested in stocks or MF in the lottery (conditional)	0.578	0.310	3675

Note. “Stocks” refer to “Stocks held directly”. Data are drawn from the June 2021 wave of the Consumer Expectations Survey (CES). Statistics are computed using sample weights.

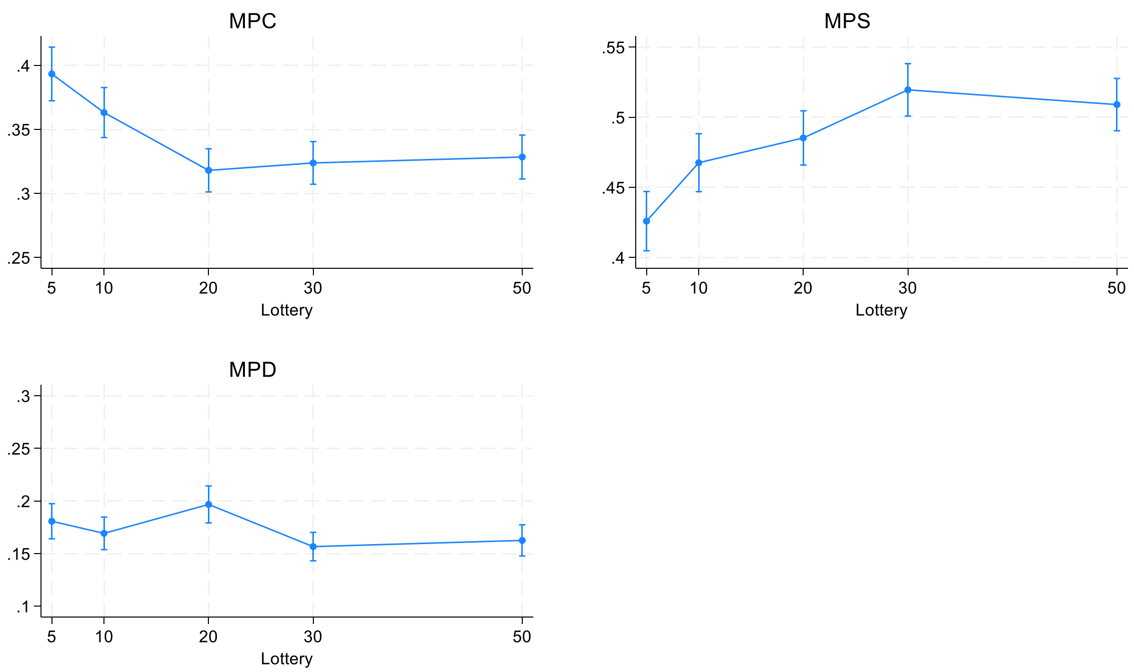

Fig. 1. Lottery prize share consumed, saved or used to reduce debt. Note. The graphs plot the averages of the Marginal Propensity to Consume (MPC), the Marginal Propensity to Save (MPS) and the Marginal Propensity to Repay Debt (MPD) by the size of the lottery prize (in thousand euro), and the associated 95 % confidence bands. Averages are computed using sample weights.

Table 2
Tobit for marginal propensity to consume, to save and to repay debt.

	MPC (1)	MPC (2)	MPS (3)	MPS (4)	MPD (5)	MPD (6)
Lottery 10k	−0.023	−0.002	0.054	0.032	−0.007	−0.008
	(0.013)*	(0.013)	(0.015)***	(0.016)**	(0.011)	(0.012)
Lottery 20k	−0.051	−0.036	0.073	0.061	0.011	0.005
	(0.012)***	(0.013)***	(0.015)***	(0.015)***	(0.011)	(0.011)
Lottery 30k	−0.049	−0.036	0.106	0.093	−0.008	−0.013
	(0.012)***	(0.012)***	(0.014)***	(0.015)***	(0.011)	(0.011)
Lottery 50k	−0.040	−0.024	0.086	0.068	0.009	0.009
	(0.012)***	(0.012)*	(0.014)***	(0.015)***	(0.011)	(0.011)
Male		−0.018		−0.006		0.032
		(0.008)**		(0.009)		(0.007)***
Age 18–34		−0.012		0.025		0.040
		(0.016)		(0.019)		(0.015)***
Age 35–49		−0.032		−0.015		0.079
		(0.016)**		(0.019)		(0.015)***
Age 50–64		0.001		−0.017		0.032
		(0.015)		(0.018)		(0.015)**
Family size		−0.001		−0.020		0.024
		(0.003)		(0.004)***		(0.003)***
High school		−0.022		−0.002		0.019
		(0.013)*		(0.016)		(0.011)*
College		−0.027		0.014		0.008
		(0.013)**		(0.015)		(0.011)
Employed		−0.015		−0.025		0.037
		(0.009)		(0.011)**		(0.009)***
Self-employed		−0.037		0.004		0.036
		(0.014)***		(0.017)		(0.014)***
Log cash-on-hand		−0.006		0.045		−0.046
		(0.003)**		(0.003)***		(0.002)***
N	9677	8559	9677	8559	9677	8559

Note. The table reports Tobit regressions for the Marginal Propensity to Consume (MPC), Marginal Propensity to Save (MPS) and Marginal Propensity to repay Debt (MPD). All regressions include country dummies. The figures are average marginal effects (unconditional expectations) and figures in parentheses are heteroskedasticity consistent standard errors. One star indicates statistical significance at the 10 %, two stars at the 5 %, three stars at the 1 %.

where $A_{it}$is the amount of stocks, W_{it} is total gross financial wealth, and α_{it} is the asset share invested in stocks. $\Delta\alpha_{it} = 0$ for those who do not invest in the stock market before and after the lottery (51 % of the sample), but can be either positive or negative for those who choose to invest in stocks (38 % and 12 %, respectively). Asset shares for stocks and mutual funds combined are defined in a similar way.[12]

4.3. Stock market participation

On average (across all lotteries) the propensity to invest in stocks is 25.7 % and the propensity to invest in stocks and mutual funds combined is 38 % (see Table 1). Actual portfolio data indicate that ownership before the lottery was 23 % and 33 %, respectively. Therefore, the wealth shocks increase direct stockholding by about 12 % of the initial prevalence (15 % including mutual funds) on average.

The upper two graphs of Fig. 2 display the proportion of those that intend to invest in risky assets for each of the five lottery gains, and the associated 95 % confidence bands.[13] The difference in participation between the smallest and the largest prize is about 10 percentage points. The other noteworthy outcome of the experiment is that even after receiving a €50,000 lottery prize, the majority of respondents would not invest in stocks or mutual funds.

Figure A2 of the Appendix shows that the positive relation between lottery gains and stock market participation applies to each of the six countries in the sample. There are also country-specific differences, as Spanish and French respondents show lower propensities to invest in stocks at any given level of the lottery prize, while in Germany and Belgium the propensity is higher. In some countries the

[12] All variables at t-1 (i.e., before the lottery gain) are predetermined, and thus their distributions are the same across all five groups due to the random assignment of lottery gains.

[13] Separate figures for mutual funds, bonds and retirement products are reported in Figure A1 of the Appendix, showing that for each of these assets there is a positive relation between the lottery prizes and ownership.

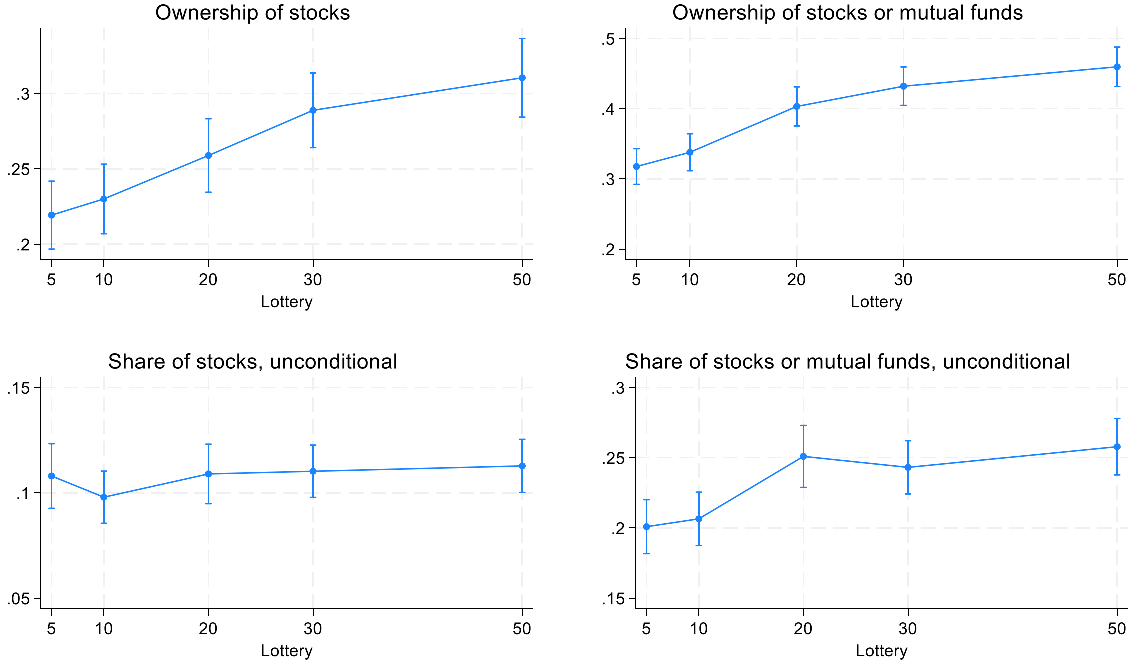

Fig. 2. Proportion of households investing and asset shares invested in stocks and mutual funds. Note. The upper two graphs plot the fraction of households that intend to invest in stocks and in stocks or mutual funds by the size of the lottery prize (in thousand euro), and the associated 95 % confidence bands. The lower two graphs plot the average asset share that households intend to invest in stocks, and in stocks or mutual funds, by the size of the lottery prize (in thousand euro), and the associated 95 % confidence bands. Asset shares are computed as the ratio of the amount invested in stocks and in stocks or mutual funds over the total amount saved. Asset shares statistics include all respondents ("unconditional" on stock market participation). Averages are computed using sample weights.

positive association between wealth and stockholding is stronger at low levels of wealth (Belgium, Italy, the Netherlands). In our regressions with the pooled sample we use country dummies to absorb this variability.[14]

The descriptive statistics are consistent with previous research using actual inheritances and lottery gains. In Denmark those who receive an unexpected inheritance equivalent to €50,000 increase stockholding by 12.9 % (Andersen and Nielsen, Table 4). In Sweden winning a lottery prize of similar range increases participation by 8.2 % (Briggs et al., p. 63). Both studies also find that most households do not enter the stock market even at high levels of wealth.[15]

It is instructive to focus on new entrants in the stock market, distinguishing between stockholders and non-stockholders before the lottery.[16] In both sub samples entry is positively correlated with the prize. For those who do not invest in stocks before the lottery, a change from the smallest to the largest prize would increase stockholding by 6 percentage points. Entry using the broader definition of stockholding is larger (9 points). As we shall see, regression analysis confirms the descriptive evidence.

4.4. Share of risky assets

In the bottom left graph of Fig. 2 we plot the part of the lottery gain invested in stocks, standardized by the total amount of the lottery gain invested in financial assets (if positive); otherwise, the asset share is set to zero. The slope of the share of stock is flat (at approximately 11 %), with no relation with the prize. For stocks and mutual funds combined, the relation increases up to €30,000 (from 20 to 25 %) and then becomes approximately linear (bottom right graph of Fig. 2). However, these statistics do not distinguish between stockholders and non-stockholders. As we shall see with regression analysis, the relation between the asset share and wealth, conditional on participation, is rather flat also for the broad definition of risky assets.

The definitions in Fig. 2 imply that participation and the asset shares are defined for all respondents, not only for those who intend to save. As a result, we base our inference on the entire sample that we use to randomly assign the lottery prizes. An alternative definition of stockholding would include only those who save positive amounts (78,1 % of the sample). The fraction of such savers is

[14] Controlling for country fixed effects is needed as also because randomization of lottery gains is done separately in each country.

[15] In addition, using a structural life-cycle model, Briggs et al. (2021) find that the model predicts much larger rates of entry following a lottery windfall, unless one is willing to assume implausibly large entry costs.

[16] Due to randomization and to the fact that pre-lottery stockholding is a predetermined magnitude, the distribution of lottery gains is the same in the two subsamples.

not random, however, and depends on the lottery amount: it is 68.5 % for the €5,000 lottery, and 82.3 % for the €50,000. This alternative definition therefore results in sub-samples that depend on endogenous selection by respondents (i.e., the save / not save decision), which in turn depends on the treatment (the lottery prizes).[17]

5. Regression analysis

We start our empirical analysis using probit regressions to estimate the decision to invest in stocks (or in stocks and mutual funds) after winning a randomly assigned prize:

$$Prob(P_{it}) = Prob\left(\gamma_0 + \sum_{j=1}^{M} \beta_j L_{jt} + \gamma_1 X_{it} + \varepsilon_{it} > 0\right) \tag{2}$$

where P_{it} is the probability that the prize is invested in stocks, L_{jt} represent the five prizes, X_{it} is a set of additional controls, and ε_{it} is an error term. Note that rather than a single lottery gain variable, we use a fully flexible specification with separate dummies for each prize.

5.1. Stock market participation

The baseline regression is reported in column (1) of Table 3. The right-hand-side variables include only four dummies for lottery gains (the €5,000 prize is the base category), and country dummies. The table reports marginal effects, with robust standard errors in parentheses. Other things equal, participation increases by 8.6 percentage points going from the lowest to the highest lottery prize.

Given the randomized nature of our treatment, even such a parsimonious specification allows us to identify genuine causal wealth effects, while adding more covariates should mainly improve the efficiency of our estimates, without affecting the causal effect of the wealth shocks. In column (2) we include age dummies, dummies for gender, education, employment, family size and log cash-on-hand, to control for resources prior to the experiment, and dummies for "trust to others" and liquidity.[18] We also include a dummy for "high financial literacy", relying on four financial literacy questions available in the CES. We take the count of correct responses as a measure of financial literacy and consider as highly literate those answering correctly three or (all) four questions (roughly 50 % of the sample).[19] Results are unaffected, and remain unchanged when we further expand the set of covariates to include expectational, preference variables or behavioral traits (investment horizon, expectations about stock market prices and their uncertainty, and a proxy for self-confidence).

Next, we focus on post-experiment entrants, by taking the first difference of Eq. (2) with respect to the pre-lottery period (t-1) and estimating the probit regression:

$$Prob(\Delta P_{it}) = Prob\left(\sum_{j=1}^{M} \eta_j L_{jt} + \Delta\varepsilon_{it} > 0\right) \tag{3}$$

where ΔP_{it} is the probability of entering the stock market, conditional on not participating before. Notice that $\Delta P_{it} = 0$ for those who participate in both periods, as well as for those who choose to never participate. All the X_{it} variables drop when differencing, as it is reasonable to assume that nothing changes in the interval in which respondents answer the lottery questions, except the lottery prize itself.[20] This implies that post-experiment stock market entry depends only on the lottery prize. Column 3 of Table 3 shows that in this specification the effect of the lottery prize is slightly lower than in the previous regressions (6.3 percentage points for the €50,000 prize), but overall the effects of the various lottery gains are similar.

Briggs et al. (2021) find similar effects of wealth shocks on stockownership among households that did not participate in equity markets before winning a (real) lottery prize. They also point out that entry costs alone (even broadly defined) cannot explain why only 8 % of households intends to invest in the stock market after winning a €50,000, unless one is willing to assume that entry costs are extremely high.[21]

Columns 4–5–6 of Table 3 show that the increase in participation in stocks or mutual funds is slightly larger than for stocks only. Going from the smallest to the largest wealth shock, the increase in participation is 12.8 percentage points (11.6 adding demographic variables to the regression). Results for entry in the stock or mutual funds markets show that the fraction of entrants increase by 8.5 percentage points for the largest prize (column 6).

[17] We check the results restricting the analysis to savers. The effects of the wealth shocks on participation and risky asset shares are qualitatively similar but a bit weaker. Due to the failure of randomization in this endogenously generated subsample, the estimated effects of wealth are unlikely to be consistent.

[18] The trust dummy equals one if the level of trust exceeds 6, on a scale from 1 to 10. The liquid dummy equals one if financial wealth is greater than one-month income. Cash-on-hand is defined as financial wealth plus current income in period t-1 (before the lottery).

[19] The wording of the trust, liquid and financial literacy questions is reported in the Appendix.

[20] Since pre-lottery stock market participation is a predetermined characteristic, treatment randomization holds also using this variable to generate subsamples or as a control. The same holds for all the subsequent samples and variables based on predetermined characteristics.

[21] They estimate median entry costs of €30,000 with a multiperiod life-cycle portfolio model, labor income risk and Epstein-Zin preferences.

Table 3
Regressions for stock market participation.

	Stocks	Stocks	Stocks First difference	Stocks or MF	Stocks or MF	Stocks or MF First difference
	(1)	(2)	(3)	(4)	(5)	(6)
Lottery 10k	0.025	0.008	0.019	0.029	0.018	0.014
	(0.014)*	(0.015)	(0.012)	(0.016)*	(0.016)	(0.013)
Lottery 20k	0.045	0.041	0.031	0.072	0.072	0.038
	(0.014)***	(0.015)***	(0.011)***	(0.015)***	(0.016)***	(0.012)***
Lottery 30k	0.083	0.075	0.058	0.120	0.109	0.071
	(0.014)***	(0.015)***	(0.011)***	(0.015)***	(0.016)***	(0.012)***
Lottery 50k	0.086	0.084	0.063	0.128	0.116	0.085
	(0.014)***	(0.015)***	(0.011)***	(0.015)***	(0.016)***	(0.012)***
Male		0.081			0.081	
		(0.009)***			(0.010)***	
Age 18–34		0.135			0.118	
		(0.019)***			(0.021)***	
Age 35–49		0.043			0.018	
		(0.019)**			(0.020)	
Age 50–64		0.001			−0.009	
		(0.019)			(0.019)	
Family size		0.006			0.006	
		(0.004)			(0.004)	
High school		−0.012			−0.015	
		(0.016)			(0.018)	
College		−0.007			0.004	
		(0.015)			(0.017)	
Employed		0.002			−0.005	
		(0.012)			(0.013)	
Self-employed		0.045			0.038	
		(0.018)**			(0.019)**	
High literacy		0.040			0.088	
		(0.010)***			(0.011)***	
High trust		0.023			0.033	
		(0.009)**			(0.010)***	
Liquidity		0.025			0.032	
		(0.013)**			(0.014)**	
Log cash-on-hand		0.023			0.044	
		(0.004)***			(0.004)***	
N	9677	8460	9677	9677	8460	9677

Note. Columns 1 and 2 report probit regressions for the probability of investing part of the prize in stocks. Column 3 reports a probit regression of the probability of entry into the stock market after receiving the prize. Columns 4 and 5b report probit regressions for the probability of investing part of the prize in stocks or mutual funds. Column 6 reports a probit regression of the probability of entry into the stock or mutual funds markets after receiving the prize. All regressions include country dummies. The figures are average marginal effects and figures in parentheses are heteroskedasticity consistent standard errors. One star indicates statistical significance at the 10 %, two stars at the 5 %, three stars at the 1 %.

The structure of our experiment allows us to check if the wealth-participation gradient varies with respondents' characteristics. We summarize graphically our main findings for the broad definition, and report in the Appendix variable definitions, regression coefficients, and results for the narrow definition of stockholding.

One recurrent argument in the literature is that information costs, above monetary costs, limit stockholding. One way to gauge the relative importance of information and transaction costs in determining the participation decision is to interact in our baseline specification (col. 1, Table 3) the lottery prizes with a dummy representing high financial sophistication. Less sophisticated investors face higher information costs, and therefore might be less sensitive to the hypothetical lottery prize, while more sophisticated investors might increase participation more at any given level of the prize.

The upper-left graph of Fig. 3 displays the results graphically from such an interaction model for the broad definition of stockholding. We plot the difference of marginal effects of the lottery prizes between groups with high and low levels of financial literacy. The relative effects are positive (in the order of 2–3 percentage points), but the 95 % confidence bands indicate that they are not statistically different from zero, regardless of the lottery prize.

Individuals with low liquid assets might be more reluctant to invest in stocks, even after receiving large windfall gains, because they might feel unprepared for emergencies or unable to borrow in case of need. Accordingly, in the upper-right graph of Fig. 3 we report differences in marginal effects when interacting lottery prizes with a dummy variable ("liquid") that equals one if households' liquid resources exceed monthly income. We find that the propensity to invest is higher only for medium prizes (between €20,000 and €30,000), although statistically insignificant at 5 %. Trust in others is associated with the level of stockholding, as shown in the baseline regressions of Table 3, but it does not affect the gradient between wealth shocks and stockholding (bottom-left graph of Fig. 3).

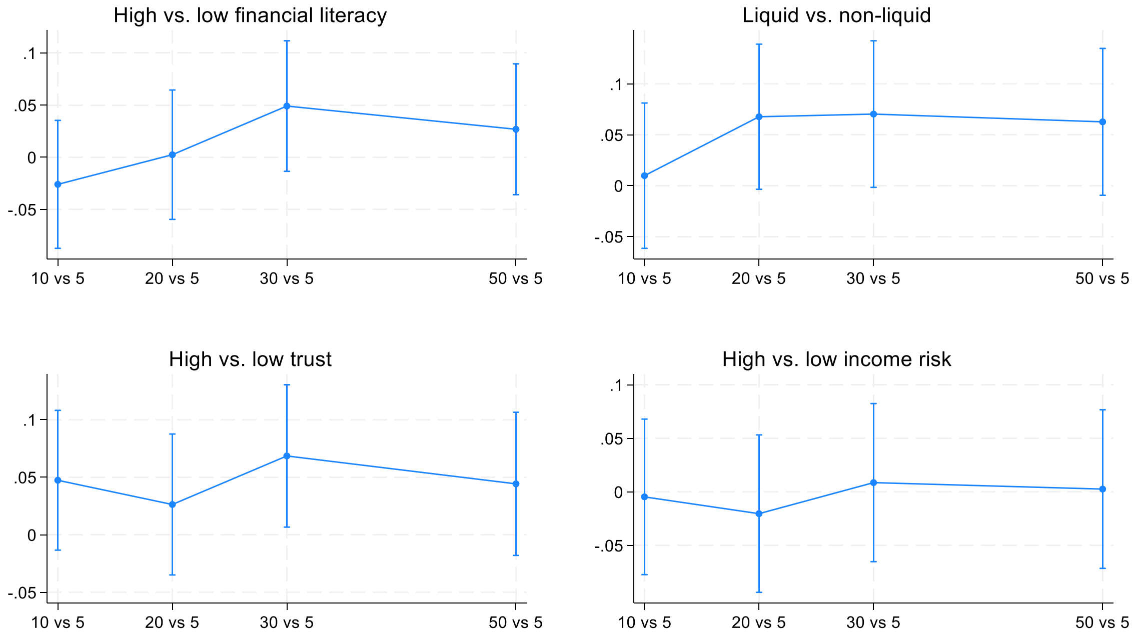

Fig. 3. Changes in the probability of investing in stocks or mutual funds, by household characteristics. Note. The graphs plot the differences of marginal effects and associated 95 % confidence intervals bands, with respect to the baseline lottery prize, for different groups of households: previous stockowners vs. non-owners (top left graph), high vs. low financial literacy (upper-right graph), liquid vs. non-liquid (bottom-left graph), high vs. low trust (bottom-left graph), high vs. low-income risk. The marginal effects are obtained from probit regressions using the same specification of column 5 of Table 3, interacting the lottery prizes with the corresponding household group dummy.

Background risk might reduce the incentive to invest in stocks, at any given level of wealth shocks. The reason is that when people face non-diversifiable risks, they might limit exposure to risks that can be avoided, such as financial risk. The bottom-right panel of Fig. 3 reports differences in marginal effects of individuals with high or low human capital risk, as proxied by the subjective standard deviation of income risk 12 months ahead. The relative effects are not statistically different from zero, suggesting that in our experiment background risk considerations do not play a role in explaining the wealth-participation gradient.

In the Appendix we check other dimensions of heterogeneity, using interactions of the prize dummies with other predetermined variables: investment horizon, expected growth of stock prices and their uncertainty, self-confidence and quartiles of cash-on-hand (or cash-on-had to income ratios). In probit regressions we find that stockholding is higher among the self-confident and optimists about the stock market, but the marginal effect differences (computed across different values of these characteristics) are not statistically different from zero (results are reported in Figures A4-A5-A6-A7 of the Appendix).

We also estimate separate regressions for the six countries of our sample. We find some heterogeneity of responses, but overall the pattern of results is similar to those of the pooled sample (results for the broad definition of the risky assets are reported in Tables A2 and A3 of the Appendix). Stockholding increases with the prize in all countries considered, and in all countries the majority of households does not invest in stocks, even at the largest lottery prize.

Overall, results suggest that wealth shocks induce an increase in stockholding between 8.4 and 12.8 percentage points for the largest prize, depending on the definition of stockholding. Wealth shocks affect also first-time entry in the stock market (between 6.3 and 8.5 points). Except for liquidity (at intermediate levels of the prizes) we find no differences in the wealth-participation gradient. Finally, the majority of respondents would not invest in stocks or mutual funds, even for the highest lottery prize.

5.2. Share of risky assets

Our randomized assignment provides an ideal setting to test the causal effect of wealth on the share of risky assets and to make indirect inferences about the relation between wealth and risk aversion, a long-standing issue in finance and economics. Consider a static model in which wealth is fully liquid and the utility function has CRRA coefficient ρ_i, the expected risk premium of investor i at time t is Er_{it}^e, and the standard deviation of the return of the risky asset is σ_{it}. In this simple framework one obtains a closed form solution of the risky asset share that is independent from wealth:

$$\alpha_{it} = \frac{1}{\rho_{ii}} \frac{Er_{it}^e}{\sigma_{it}^2}. \tag{4}$$

In other contexts, investors exhibit Decreasing Relative Risk Aversion (DRRA), arising for instance from habits, subsistence consumption or committed expenditures. Let's express the optimal share invested in stocks as:

$$\alpha_{it} = h\left[W_{it},\ Er^{e}_{it},\ \sigma^{2}_{it},\ f_{it}\right] \tag{5}$$

where f_{it} includes various observed and unobserved variables, including demographic characteristics, transaction costs, the level and riskiness of human capital, habits and other preference traits that affect the optimal risky asset share, and are potentially varying over time.

Eq. (5) suggests that there are many challenges in estimating the relation between the risky asset share and wealth with observational data. OLS estimates likely lead to omitted variable and endogeneity bias, for instance because households may have become richer because they have invested in the risky asset. The crucial advantage of our experiment is that wealth shocks are exogenous and randomly distributed by construction, and thus we can cleanly estimate their causal effect on risky asset shares. Moreover, we can safely assume that all other characteristics (preferences, human capital, income risk, etc.) are constant in the time interval in which respondents answer the hypothetical lottery questions. In other words, in our experiment the only reason why the asset share can change is the random shock.

Our approach is to consider a linear approximation of Eq. (5), and to estimate a regression relating the risky asset share to lottery gains of different size:

$$\alpha_{it} = \gamma_0 + \sum_{j=1}^{M} \pi_j L_{jt} + \gamma_1 X_{it} + \mu_{it} \tag{6}$$

where $j = 1,\ldots,M$ indicates the increasing levels of the hypothetical gains L, the π_j coefficients measure their effects on asset shares, and the term μ is measurement error. The sign and magnitude of π_j speaks to risk aversion. A finding that the risky asset share is independent from wealth ($\pi_j = 0$) would be consistent with CRRA, while positive and increasing values of π_j would support models with DRRA (increasing aversion with negative and decreasing values).

Estimating Eq. (6) with our data shares many of the advantages discussed for Eq. (5). First, we measure the portfolio response to unanticipated and exogenous wealth shocks of different size randomly distributed in the sample. Second, given our randomization, the estimated π_j measures the causal effect of wealth shocks on the risky asset share, as any omitted variables are by design uncorrelated with our treatment variable, that is $cov\left(L_{jt},\ \mu_{it}\right) = 0$. Therefore, the omission of these variables from Eq. (6) should not affect the estimated π_j. Third, measurement error in wealth shocks is ruled out by our experimental design, which fixes the euro amounts of the lottery prizes. In short, our experiment establishes a causal link between the risky asset share and change in wealth, holding all other variables constant, including changes in preferences, demographics, and human capital.

Given that the asset share cannot be negative, in Table 4 we report marginal effects from Tobit regressions. In column (1) the asset share increases from 0.4 % (for the €10,000 prize) to 1.5 % (for the €50,000 prize), relative to baseline. The effects are precisely estimated for prizes of €20,000 and above, and unchanged when adding demographic variables, cash-on-hand, and dummies for literacy, trust, and liquidity (column 2).

As we did with the probits for participation, in column (3) of Table 4 we estimate by OLS the first difference of Eq. (6):

$$\Delta\alpha_{it} = \sum_{j=1}^{M} \pi_j L_{jt} + \Delta\mu_{it} \tag{7}$$

To compute the change in the asset share, we make use of information on the pre-existing, and thus pre-determined, level of wealth.[22] In column (3) the coefficients of the lottery dummies are slightly negative, and significant only for the €20,000 and €50,000 prizes. Notice however that in this specification we exclude all those who don't invest in stocks before the lottery, so the number of observations drops from 9677 to 2191.

In columns 4–5–6 of Table 4 we repeat the estimation using the broader definition of risky assets. Columns (4) and (5) show that the asset share is rather insensitive to wealth shocks for relatively small shocks. For prizes of €30,000 and above, the asset share increases by 1.7 % with respect to the pre-lottery portfolio allocation. Column (6) shows that for the change in the broad definition of the asset share there is a small positive coefficient only for the largest prize (1.7 %).

Next, we explore potential sources of differences in the gradient between wealth shocks and changes in the risky asset share (broad definition). The graphs are obtained by interacting the prize dummies with the same set of variables used in Section 5.1 for participation (financial literacy, liquidity, trust, income risk), and plotting the differences in marginal effects relative to the baseline prize. Fig. 4 shows that the effect of wealth shocks on risky asset shares hardly changes across these different groups. The relative effects of financial literacy increase with the prize, but are not statistically different from zero. Only in the upper-right graph the risky asset share of high liquid respondents exhibits a slightly higher sensitivity to lottery prizes of €20,000 and €30,000 (about 6 %), relative to low liquid ones. Figures A8-A9-A10 in the Appendix report that results are robust using different definitions of stockholding and interaction models capturing potential differences due to the investment horizon, expected growth of stock prices and their uncertainty, and self-confidence.

[22] Since the change in this risky asset share can be negative, using OLS is appropriate.

Table 4
Regressions for the share of risky assets.

	Stocks (1)	Stocks (2)	Stocks, holders only First difference (3)	Stocks or MF (4)	Stocks or MF (5)	Stocks or MF, holders only First difference (6)
Lottery 10k	0.004	−0.002	−0.002	0.006	0.001	−0.000
	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.005)
Lottery 20k	0.008	0.007	−0.019	0.015	0.011	−0.000
	(0.004)**	(0.004)	(0.007)***	(0.004)***	(0.004)***	(0.006)
Lottery 30k	0.016	0.013	−0.004	0.025	0.017	0.009
	(0.004)***	(0.004)***	(0.007)	(0.004)***	(0.004)***	(0.007)
Lottery 50k	0.015	0.014	−0.013	0.024	0.016	0.017
	(0.004)***	(0.004)***	(0.007)*	(0.004)***	(0.004)***	(0.007)**
Male		0.024			0.020	
		(0.003)***			(0.002)***	
Age 18–34		0.032			0.015	
		(0.006)***			(0.005)***	
Age 35–49		0.009			−0.003	
		(0.006)			(0.005)	
Age 50–64		−0.001			−0.006	
		(0.006)			(0.005)	
Family size		0.001			0.001	
		(0.001)			(0.001)	
High school		−0.001			−0.001	
		(0.005)			(0.004)	
College		0.000			0.003	
		(0.004)			(0.004)	
Employed		0.001			0.000	
		(0.003)			(0.003)	
Self-employed		0.015			0.009	
		(0.005)***			(0.004)**	
High literacy		0.015			0.024	
		(0.003)***			(0.002)***	
High trust		0.005			0.007	
		(0.003)**			(0.002)***	
Liquidity		0.007			0.006	
		(0.004)*			(0.003)*	
Log cash-on-hand		0.009			0.012	
		(0.001)***			(0.001)***	
N	9677	8460	2191	9677	8460	3122

Note. Columns 1 and 2 report Tobit regressions and the dependent variable is the asset share of the prize invested in stocks. Column 3 reports OLS estimates from a sample that includes only stockholders before the experiment, and the dependent variable is the change in the asset share invested in stocks. All regressions include country dummies. Columns 4 and 5 report Tobit regressions and the dependent variable is the asset share of the prize invested in stocks or mutual funds. Column 6 reports OLS estimates from a sample that includes only holders of stocks or mutual funds before the experiment, and the dependent variable is the change in the asset share invested in stocks or mutual funds. The figures in columns 1, 2, 4 and 5 are average marginal effects (unconditional expectations) and figures in parentheses are heteroskedasticity consistent standard errors. All regressions include country dummies. One star indicates statistical significance at the 10 %, two stars at the 5 %, three stars at the 1 %.

To summarize, we find that the asset share is rather insensitive to wealth, using different definitions of stockholding and estimation approaches. Differences across groups are not significant, except when considering liquid consumers confronted with intermediate prizes. One interpretation of the survey experiment is therefore that our results are consistent with the hypothesis of CRRA $\left(\pi_j = 0\right)$, except when wealth shocks are very large (€30,000 and above), in which case risk aversion is slightly decreasing.

One possibility is that our findings may be, at least partly, a reflection of the fact that the wealth effect might vary with wealth itself, or that the size of the wealth shocks that we consider is not large enough. To address the first issue, in Fig. 5 we check if the relation between the risky asset share and wealth shocks depends on the position of the individual in the wealth distribution. We estimate the relation for each wealth quartile of cash-on-hand, and for each quartile of the ratio between cash-on-hand and income. In no case we find significant differences across quartiles.

To see why the second issue is relevant, take a case where utility has a habit term given by a minimum level of wealth $\overline{W}$ and the utility function is $u = \frac{(W-\overline{W})^{(1-\gamma)}}{1-\gamma}$. Relative risk aversion is $RRA = \gamma \frac{1}{\left(1-\frac{\overline{W}}{W}\right)}$ and is decreasing in W. It follows that the sensitivity of RRA to changes in wealth is low for rich individuals with low $\left(\frac{\overline{W}}{W}\right)$. Instead, adding wealth to the portfolio of a poor individual with high

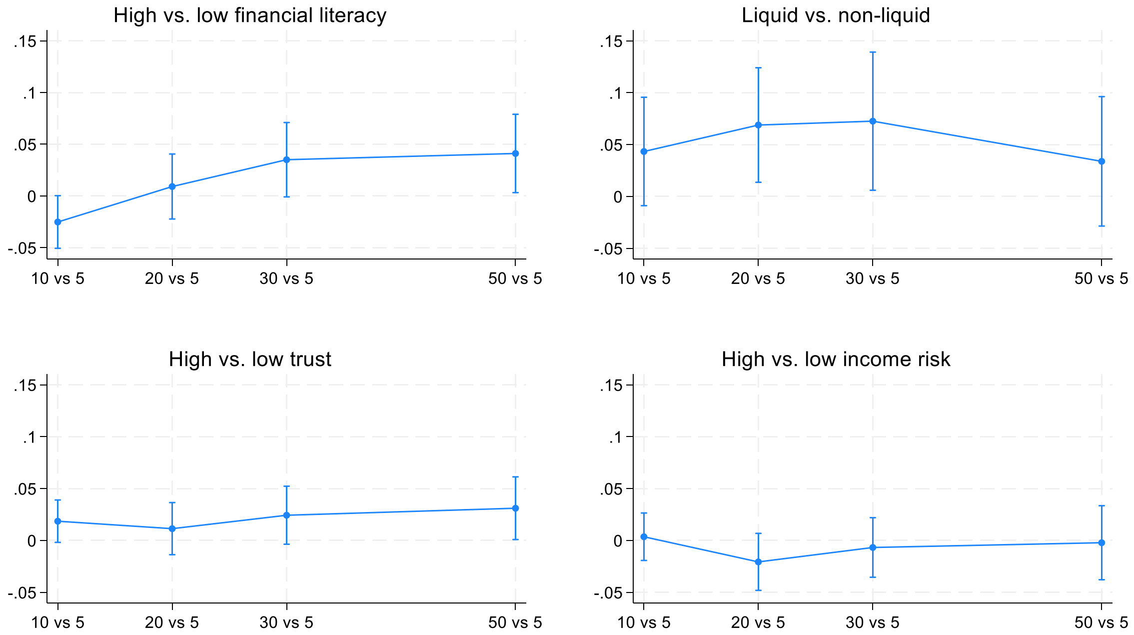

Fig. 4. Change in the asset share invested in stocks or mutual funds, by households' characteristics. Note. The graphs plot the differences of marginal effects and associated 95 % confidence bands, with the respect to the baseline lottery prize, for different groups of households: high vs. low financial literacy (top-left graph), liquid vs. non-liquid (top-right graph), high vs. low trust (bottom-left graph), high vs. low-income risk (bottom-right graph). The estimated effects are obtained using the same specification of column 5 of Table 4, interacting the lottery prizes with the corresponding household group dummy.

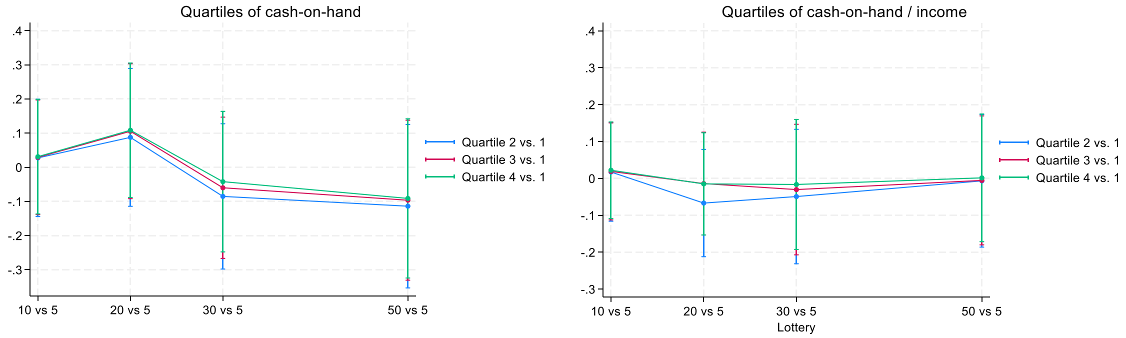

Fig. 5. Change in the asset share invested in stocks or mutual funds, by cash-on-hand quartiles. Note. The graphs plot the differences of marginal effects and associated 95 % confidence bands with respect to the baseline lottery prize, for different cash-on-hand quartiles (left-graph) and different ratios of cash-on-hand to income (right graph). In all cases, the first quartile is the comparison group. The estimates are obtained from OLS regressions interacting the lottery prizes with the quartile dummies.

$\left(\frac{\overline{W}}{W}\right)$ has a large effect on RRA and portfolio allocation. In the experiment, those who receive a high prize starting from a position close to $\overline{W}$, that is, those with high $\left(\frac{\overline{W}}{W}\right)$, will be the most sensitive.

One way to check for the presence of the habit term is to split the sample by wealth before the lottery (a predetermined characteristic, and thus orthogonal to the lottery), conditioning on those who already own stocks before the lottery. In this conditional sample of stockholders, the most interesting groups are those with low wealth (say, €5,000 or €10,000) who receive large lotteries. We don't find significant effects of prizes, but in doing this we need to focus on stockholders with low actual wealth. However, since there are relatively few such observations, the coefficients are not precisely estimated.[23]

[23] For instance, the group investing in stocks or mutual funds before the lottery, receiving a €50,000 prize and having less than median cash on hand (€18,900) includes only 103 observations. These results are available on request. We thank Luigi Guiso for suggesting this test.

Table 5
Regressions with actual portfolio data.

	Participation		Risky Asset Share	
	Stocks (1)	Stocks or MF (2)	Stocks (3)	Stocks or MF (4)
Lottery 10k	0.024	0.033	0.007	0.011
	(0.001)***	(0.001)***	(0.000)***	(0.000)***
Lottery 20k	0.058	0.077	0.018	0.026
	(0.001)***	(0.003)***	(0.001)***	(0.037)***
Lottery 30k	0.083	0.108	0.025	0.037
	(0.004)***	(0.004)***	(0.001)***	(0.001)***
Lottery 50k	0.120	0.153	0.036	0.054
	(0.005)***	(0.005)***	(0.002)***	(0.002)***

Note. Regressions use actual portfolio data from the June 2021 Consumer Expectations Survey (CES). In columns 1 and 2 the dependent variable is actual participation in stocks or mutual funds. In columns 3 and 4 the left-hand-side variable is the actual asset share invested in stocks or mutual funds (before receiving the hypothetical lottery prize). To ensure comparability with Tables 3 and 4, the table shows the marginal effect of increasing observed gross financial assets from 5k to the amount of lottery gains shown on the probability of pre-experiment stock ownership and on the conditional pre-experiment risky financial asset share. We compute, using the corresponding regression equation, the predicted outcomes in both cases, and then take the difference in these predicted outcomes. The average difference, and associated standard error, can be interpreted as the average differential effect of increasing wealth by each of the lottery prize, compared to 5k prize. The regressions models used are probits for participation and Tobits for the risky asset share. The control variables included correspond to the full specification shown in columns (2) and (5) of Tables 3 and 4. One star indicates statistical significance at the 10 %, two stars at the 5 %, three stars at the 1 %.

6. Actual vs. hypothetical data

In this section, we compare the estimated responses to hypothetical lottery prizes with the responses estimated with actual data on stockholding and risky asset shares also collected in June 2021 special-purpose survey. Our aim is to generate, using data on actual financial wealth, results that can be compared as much as possible to our baseline results shown in Tables 3 and 4 (i.e., the effects generated through our experiment by comparing a wealth gain of €10,000, €20,000, €30,000 and €50,000 to the baseline of €5,000).

The estimates shown in Table 5, refer to four differential effects in wealth increments using regressions in which actual financial wealth is the variable of interest. To achieve this, we regress participation and asset shares on the log of financial wealth and the same control variables as in Tables 3 and 4.[24] Then, taking as an example the reported effect for the €10,000 prize, we increment for each household financial wealth by €5,000 and €10,000 and compute the predicted outcomes in both cases. Last, we calculate and report the average difference in these predicted outcomes. In a similar way, for comparison with the results in Tables 3 and 4, we compute the reported effects for the other €20,000, €30,000 and €50,000 prizes.[25]

We find that participation and asset shares are more strongly correlated with wealth in actual data than in the survey experiment. For instance, for participation, the increase in stockholding is 12 percentage points against 8.4 percentage points in our experiment for the largest prize (column 2 of Table 3), and 15 percentage points against 11.6 percentage points using the broad definition. Also, in actual data asset shares of stocks increase by 3.6 % for the largest prize (5.4 % for the broad definition), while in the experiment the corresponding increase is about 1.5 % (columns 2 and 5 of Table 4). A likely explanation for this difference is that in actual data wealth is correlated with other unobserved confounding factors that are associated with stockholding and risky asset shares, while the random assignment of lottery prizes allows estimating the pure effect of wealth shocks. For instance, it is likely that investors who are more financially sophisticated, better-informed and less risk averse tend to have higher wealth, leading to a stronger correlation between wealth and stockholding in actual data.

7. Implications for macro models

Our finding that the risky asset share is rather insensitive to wealth shocks is consistent with the notion that preferences are characterized by CRRA in the relevant range of our experiment. This is in line with the widely cited paper of Chiappori and Paiella (2011), who find in panel data that changes in asset shares are uncorrelated with changes in wealth. Notice that even if we conclude that wealth shocks don't affect the risky asset share on average (also within subsamples defined by characteristics), one cannot identify the average CRRA parameter, without making additional and strong assumptions.

To see this point, let's consider the standard portfolio model of Eq. (4). Suppose that all individuals have CRRA utilities with coefficients equal to ρ_i, and the same expectations of returns Er^e and volatility σ^2. Then one could infer the cross-sectional distribution

[24] As with our results from the experiment, the regressions models used were probits for participation and Tobits for the risky asset share. The control variables included corresponded to the full specification shown in columns (2) and (5) of Tables 3 and 4.

[25] Given that these discrete differences are heterogeneous across households, we computed the average difference across all households. This average difference, which is recorded as the €10,000 coefficient in Table 5, can thus be interpreted as the average differential effect of increasing financial wealth by €10,000 compared to €5,000, which is what the coefficient of the dummy variable representing the €10,000 prize denotes in Tables 3 and 4. We repeated this process for each of the other prizes to calculate the respective average discrete difference.

of CRRA from the relation: $\rho_i = \frac{1}{a_i} \frac{Er^e}{\sigma^2}$. Thus, even in the simplest framework, the distribution of CRRA is identified only up to the scale factor $\frac{Er^e}{\sigma^2}$. If one further assumes, as in Chiappori and Paiella (2011), that the expected excess return and the volatility of stocks are calibrated to $r^e = 0.04$ and $\sigma^2 = 0.04$, one can recover from the data the cross-sectional distribution of CRRA (up to a scale factor, assumed to be 1 in our example).

In Figure A11 of the Appendix we plot the cross-sectional distribution of CRRA using data on the actual risky asset share (stocks and mutual funds), eliminating observations with risky asset shares below 5 %. The median of the CRRA distribution is 2.6 and the mean is 3.9, with a standard deviation of 3.5. The figure is for illustrative purpose only, and it is meant to show that our estimates don't imply that investors have the same CRRA parameter. Rather, they are consistent with considerable heterogeneity in the CRRA parameters, which even in the simplest portfolio model are identified up to a scale factor.

These findings might be useful when evaluating the ability of Heterogeneous Agents New Keynesian (HANK) models to replicate particular features of the wealth and consumption distributions. HANK models consider various sources of heterogeneity, such as the composition of wealth and returns between liquid and illiquid assets, access to credit, and income risk. Recent papers suggest that heterogeneity in preferences (discount rates, elasticities of intertemporal substitution, and relative risk aversion) should also be considered.

Cozzi (2013) shows that heterogeneity in risk aversion - calibrated using the Chiappori and Paiella (2011) estimates - can match the observed degree of wealth inequality in the U.S., except for the top of the wealth distribution. Carroll et al. (2017) find that heterogeneity in impatience affects the size and the heterogeneity of the MPC. Aguiar et al. (2023) suggest that heterogeneity in the intertemporal elasticity of substitution and in discount factors can explain several features of the consumption distribution in the PSID. Kaplan and Violante (2022) conduct a systematic analysis of various sources of heterogeneity and conclude that preference heterogeneity is one ingredient that helps in producing realistic MPC distributions, reducing the need to assume a large gap in the returns to liquid and illiquid wealth in their two-assets model.

By providing evidence in support of CRRA heterogeneity, our paper suggests that it would be interesting to explore further to what extent this HANK models with heterogenous preferences are able to explain features of the wealth and consumption distributions that are not easily captured by differences in income risk, borrowing constraints or asset liquidity. An important implication of our results for HANK models is that preference heterogeneity should be of a form that does not lead to the risky asset share depending on wealth shocks on average, with the possible exception of very large shocks.

8. Summary

We implement a novel set of questions in a population representative survey to elicit household-specific propensities to invest and consume out of positive wealth shocks through a randomized assignment of hypothetical lottery gains ranging from €5,000 to €50,000. This set-up allows us to estimate the causal effect of wealth shocks on risky asset ownership and conditional asset shares and to shed new light on investors' risk preferences. We find that stock market participation is positively related to wealth shocks, and that wealth shocks favor first-time entry in the stock market. The effect of wealth on participation is not large, however. Even after receiving a wealth windfall gain of €50,000, participation of non-stockholders only increases by between 8.4 and 12.8 percentage points. And even after receiving such a large prize, over 60 % of respondents would not invest in stocks, suggesting that transaction costs are only partly responsible for limited stock market participation. Rather, for many potential investors, reasons for non-participation likely also include information costs as well as beliefs about stock prices, lack of trust, inertia, and other behavioral biases. The survey uncovers limited differences of responses in relation to financial sophistication, liquidity, trust, income risk and other variables.

The asset share invested in stocks (or stocks and mutual funds), conditional on participation, is rather insensitive to the size of the wealth shocks, with a quantitatively small exception for the largest prizes. Hence, this finding is clearly observed for a range of wealth shocks that are realistic for most households. Comparison of data from our experiment and actual data on wealth and stockholding suggest that regressions based on actual data tend to overestimate the effect of wealth shocks on the risky asset share. Our results are consistent with the notion that preferences are on average well characterized by CRRA, and also point to some important heterogeneity in the level of risk aversion across individuals. Thus, it would be interesting to explore further to what extent structural models with heterogenous preferences are able to explain features of the wealth and consumption distributions that are not easily captured by differences in income risk, borrowing constraints or asset liquidity.

Data availability

Data will be made available on request.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2024.103632.

References

Aguiar, M., Bils, M., Boar, C., 2023. Who are the hand-to-mouth?. In: NBER Working Paper 26643.
Andersen, S., Nielsen, K.M., 2011. Participation constraints in the stock market: evidence from unexpected inheritance due to sudden death. Rev. Financ. Stud. 24, 1667–1697.
Armantier, O., Topa, G., van der Klaauw, W., Zafar, B., 2016. An overview of the survey of consumer expectations. In: Federal Reserve Bank of New York Staff Reports No. 800.
Bachmann, R., Topa, G., van der Klaauw, W., 2022. Handbook of Economic Expectations. Elsevier, Amsterdam.
Bilias, Y., Georgarakos, D., Haliassos, M., 2010. Portfolio inertia and stock market fluctuations. J. Money, Credit Banking 42, 715–742.
Briggs, J., Cesarini, D., Lindqvist, E., Östling, R., 2021. Windfall gains and stock market participation. J. Financ. Econ. 139, 57–83.
Brunnermeier, M., Nagel, K.S., 2008. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals. Am. Econ. Rev. 98, 713–736.
Calvet, L.E., Campbell, J.Y., Sodini, P., 2009. Fight or flight? Portfolio rebalancing by individual investors. Q. J. Econ. 124, 301–348.
Calvet, L., Sodini, P., 2014. Twin picks: disentangling the determinants of risk-taking in household portfolios. J. Finance 69, 867–906.
Carroll, C., Slacalek, J., Tokuoka, K., White, M., 2017. The distribution of wealth and the marginal propensity to consume. Quant. Econom. 8, 977–1020.
Chiappori, P.A., Paiella, M., 2011. Relative risk aversion is constant: evidence from panel data. J. Eur. Econ. Assoc. 9, 1021–1052.
Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., van Rooij, M., 2019. Asymmetric consumption effects of transitory income shocks. Econ. J. 129, 2322–2341.
Cocco, J., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. Rev. Financ. Stud. 18, 491–533.
Cozzi, M., 2013., Risk aversion heterogeneity, risky jobs and wealth inequality. Mimeo, Queen's University.
Crawley, E., Theloudis, A., 2024. Income shocks and their transmission into consumption. Forthcoming in the Encyclopedia of Consumption.
ECB, 2021. ECB Consumer Expectations Survey: an overview and first evaluation. In: Occasional Paper Series, No. 287, December.
Gelman, M., 2021. What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics. J. Monet. Econ. 117, 521–542.
Georgarakos, D., Kenny, G., 2022. Household spending and fiscal support during the COVID-19 pandemic: insights from a new consumer survey. J. Monet. Econ. 129, S1–S14.
Guiso, L., Sodini, P., 2013. Household finance: an emerging field. Handbook of the Economics of Finance. Elsevier, Amsterdam.
Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. J. Finance 63, 2557–2600.
Havranek, T., Sokolova, A., 2020. Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not. Rev. Econ. Dyn. 35, 97–122.
Jappelli, T., Pistaferri, L., 2014. Fiscal policy and MPC heterogeneity. Am. Econ. J.: Macroecon. 6, 107–136.
Kaplan, G., Violante, G.L., 2022. The marginal propensity to consume in heterogeneous agent models. Annu Rev. Econom. 14, 747–775.
Merton, R., 1971. Optimum consumption and portfolio rules in a continuous-time model. J. Econ. Theory. 413, 373–413.
Meeuwis, M., 2022. Portfolios Wealth Fluctuations and Risk preferences: Evidence from U.S. investor. Mimeo. Washington University in St. Louis.
Paravisini, D., Rappoport, V., Ravina, E., 2017. Risk aversion and wealth: evidence from person-to-person lending portfolios. Manage Sci. 63, 279–297.
Shapiro, M.D., Slemrod, J., 2003. Did the 2001 tax rebate stimulate spending? Evidence from taxpayer surveys. In: Poterba, J.M. (Ed.), Tax Policy and the Economy 17. MIT Press, Cambridge, pp. 83–109.
Stantcheva, S., 2023. How to run surveys: a guide to creating your own identifying variation and revealing the invisible. Annu Rev. Econom. 15, 205–234.
van Rooij, M., Lusardi, A., Alessie, R., 2011. Financial literacy and stock market participation. J. Financ. Econ. 101, 449–472.
Wachter, J., Yogo, A.M., 2010. Why do household portfolio shares rise in wealth?. In: NBER Working Paper No. 16316.