

Wealth Accumulation: The Role of Others

Michael Haliassos

March 2024

Prepared for presentation at the
Pension Research Council Symposium, May 2-3, 2024
'Household Retirement Saving, Investment, and Spending:
New Lessons from Behavioral Research'

Wealth Accumulation: The Role of Others

Abstract

This chapter presents findings from four recent projects on the effects of neighbors and other peers, financial advisors, and exogenous stressors on household wealth accumulation. First, a larger proportion of neighbors with college-level economics or business education tends to promote retirement saving. Second, greater local wealth inequality, combined with wealth mobility, tends to motivate college-educated households at the beginning of their economic lives to take more asset risks later in life and thus achieve greater wealth, leaving the less educated behind. Third, financial advice from unbiased professionals and from peers or elders is affected differently by the characteristics of both the advisor and the advisee. Given the differential access of household groups to professional advice, this has implications for accumulation behavior and stock market participation across the wealth distribution. Finally, background stressors, such as crises and wars but also personal problems, occupy people's minds as they make saving decisions. In an incentivized online experiment, background stress imposing cognitive load consistently dampened consumption and promoted saving.

Keywords:

JEL Codes:

Michael Haliassos

Goethe University Frankfurt and CEPR
House of Finance, Goethe University Frankfurt
Theodor-W.-Adorno-Platz 3, PF H32, 60323
Frankfurt am Main, Germany
+49 69 798 33804.

Haliassos@wiwi.uni-frankfurt.de

Following the pioneering work of Manski (1993), the importance of peer effects on individual financial behavior has been extensively scrutinized. Peer effects have recently been part of an active research agenda in the rapidly expanding area of household finance, surveyed in Gomes et al. (2021). This chapter presents findings from four recent projects in which I have collaborated with several coauthors on the effects of family members, neighbors, other peers, professional advisors, and external stressors on household financial behavior, as it relates to wealth accumulation and preparation for retirement.

The first project focuses on what Manski (1993) called ‘exogenous peer effects,’ namely those arising from characteristics of peers, but not directly by their behavior. The characteristic under study is the financial literacy of neighbors with whom individuals interact, a phenomenon that we have studied under the term ‘financial literacy externalities’ (Haliassos et al. 2020). The second project turns to an analysis of the effects of being exposed to different levels of local wealth inequality at the start of one’s career, for subsequent risk taking, wealth accumulation, and rank in the cohort’s wealth distribution (Haliassos et al., 2023). The third project uses a German sample of professional financial advisors and laypersons with at least basic understanding of financial matters to compare the advice provided by both for exogenously-generated vignettes of potential investors (Rumpf et al, 2024). The idea is to compare the range of advice given by professional advisors and peers, taking into account both observable and unobservable advisor heterogeneity. While the three first projects focus on information, knowledge, and inspiration likely to be provided by peers, the fourth project focuses on the cognitive load imposed by others (‘exogenous stressors’), and on whether and how this interferes with important financial decisions of the household (Assenza et al., 2023). Such load, which burdens and distracts households while they are trying to make important choices, can

arise in the context of aggregate crises, such as a pandemic or war, or personal shocks such as life events, stressful occupational environments, or persistent health issues. The last study involves an online survey and experiment, conducted with a sample representative of the French population. In what follows, we provide an overview of the interesting research questions, our methodology, and our findings.

Neighbors' Financial Literacy

Unlike most previous research on financial literacy which focuses on the effects of own financial literacy on financial behavior, Haliassos et al. (2020) showed that having financially literate neighbors promoted people's own participation in retirement accounts and stockholding over both the medium and longer runs. This occurred by conveying their knowledge of matters related to economics and business, provided that the recipient of information could understand it.

Yet not surprisingly, such an undertaking faces many econometric and conceptual challenges. First, it is difficult to secure financial data indicating a household's precise location, as well as the key characteristics of its neighbors. Indeed, confidentiality concerns often delete location from publicly-available datasets, to prevent disclosure of households' identity. Second, even if one can obtain such information, locational proximity does not establish an exogenous influence of neighbors on any given household's behavior. For instance, a household may have chosen to live in that neighborhood because it shared common preferences or occupations with those neighbors. Even if the neighbors observably participated in retirement accounts or in stockholding, common portfolio behavior could reflect common preferences or common shocks to which neighboring households were also exposed. Further, existing research indicates that

few people discuss their personal financial matters with neighbors, even when they have large social circles.¹ Moreover, confidentiality concerns prevent allow researchers from eavesdropping on conversations and interactions between households and their neighbors, even if the most advanced dataset on friends' networks provide information on digital social networks. Therefore, an outstanding question is whether the interactions involve learning, or rather, mere imitation of neighbors' portfolios. Fourth, better-educated people can more easily collect and process information, so it is not clear that having a business or economics education is what matters for neighbors' beneficial effects. Finally, it is of interest to determine whether living near suitably educated neighbors suffices for to generate financial literacy externalities, or whether it also requires recipient households to be educated so they can process and apply the information obtained.

Our project (Haliassos et al. 2020) solved many of these problems by drawing on an unusual natural experiment implemented in Sweden 1987-1991, in 277 of 284 municipalities. Informed by the high quality of Swedish administrative data collected for purposes of taxing wealth during 1999-2007, and the generally high level of detail and tracking that Swedish data policies allow, we study a natural experiment where the Swedish government allocated destitute refugees arriving in Sweden to apartments that randomly became available. This was introduced to counter the typical tendency of refugees to locate in big cities, contributing to congestion and associated problems. This is as close as one probably gets in a free society to randomly assign allocation of people to areas where they are confronted with different configurations of peers.

It is worth describing the key components of this data configuration, especially because of the immense obstacles currently placed on researchers by the difficulty of governments to distinguish between respecting confidentiality of personal data and making such data

anonymously available to researchers. To confirm that the allocation neighborhoods were random, the authorities did consider their education, marital status, and neighbors speaking their language in the broader area were considered. Importantly, researchers could observe and control for these refugee characteristics, as well as for any other information available to the authorities, as no interviews were conducted. While refugees were asked to express preferences, there was no indication that the authorities allocated refugees according to their preferences. Also, the dataset allows researchers to track refugees' location over the following 20 years and observe how they 'voted with their feet.' By 1999, about three quarters of the refugees had relocated from their initial points of placement, yet such relocation took time, as destitute refugees had to first set themselves up, get jobs, and accumulate the means to relocate. Refugees spent 5.4 years in their initially assigned parish on average, and 8.7 years in the broader area of the initial municipality.

Our focus on destitute refugees helps us avoid endogenous sorting of refugees into areas with specific characteristics, yet it offers no hope of finding interesting household portfolio behavior at the time of their placements. If researchers are to be able to understand any longer-term effects of such initial placement on the tendency to save for retirement or to take financial risk, they require detailed portfolio data 10 to 20 years later, and an ability to link the refugee placement data to the portfolio data. Fortunately, detailed portfolio data became available in Sweden between 1999 and 2007;² not from the individuals themselves, but their financial and other relevant institutions. All taxable assets were covered included debts written off the assets for tax purposes. Data accuracy was carefully monitored by the tax authorities, and there were heavy penalties for misrepresentations leading to tax evasion. Such invaluable data allow

researchers to study the ‘long shadow’ of initial placement features on subsequent financial behavior.

The central finding of our study was that those placed in areas with higher shares of college-educated neighbors who studied business and economics were more likely to participate in private retirement accounts 10-15 years later, and holding stocks 15-20 years later. Strikingly, these effects were evident only for those who had at least high school certificates. Indeed, this is what we would expect if such participation were not an act of pure imitation, but instead if it required the processing of information and knowledge received from neighbors. Of course, these effects may have arisen from exposure to the neighbors themselves, or to other features of the neighborhood. After all, one would expect financially knowledgeable people to live in areas with adequate financial infrastructure such as banks, insurance companies, and financial advisors catering to their interests and preferences.³ For this reason, the study included specific controls for the electoral district in which each refugee’s apartment was located, which incorporates fixed factors present in that location and relevant for the two asset participation outcomes. The effects were estimated controlling for such common environments.

To support the idea that neighbors’ business and economics education matters, the Swedish data permit us to consider the share of neighbors with quantitative education more generally, along with those college educated (regardless of content). Our estimates produce much smaller portfolio effects of the former, and no effects of the latter type of neighbors. This is consistent with the idea that neighbors who convey relevant content are more effective, not simply those with the ability to collect, understand, and convey general information.

Further testing shows that the share of neighbors participating in the corresponding asset had a smaller effect on the refugees’ subsequent participation compared to that of

knowledgeable neighbors; also, the initial share of neighbors with business and economics education who *did not hold* the financial asset in question still had a significant effect. This was true even when the share of holders was controlled in the model. Indeed, the sharper findings are consistent with the requirement of a high school certificate for the refugees, and of the content of education for the neighbors. Specifically, if imitation were the main mechanism driving the participation effect in private retirement accounts and stockholding, refugees would not need to understand and process information, and neighbors would not need to convey their approach to retirement and financial risk taking.

Lacking social network data on interactions with peers, it remains to be determined whether one might argue that refugees were interacting with their neighbors. Our study varied factors affecting the probability of interaction with neighbors, to look for significant changes in the estimated magnitude of peer effects. We showed that peer effects were operative in areas where Swedish neighbors were more positively predisposed to immigrants, and where there was a critical mass of knowledgeable neighbors. Effects were also stronger for refugees who had children at the time of entry, and thus greater incentives and opportunities to interact with other parents through schools. These findings are consistent with interactions producing these effects on retirement and financial risk taking.

While consistency across these results is encouraging for researchers, it may also concern those worried about distributional implications of peer effects. That is, financial literacy externalities generate a social multiplier for attempts to provide financial education to some, which raises the chances of financial education programs to be economically viable. Yet we cannot rely heavily on such externalities to improve the financial behavior of disadvantaged socioeconomic groups. The tendency of people to associate with others like them, termed

‘homophily’, means that low socioeconomic groups are both unlikely to be interacting with knowledgeable others, as well as unable to interpret the information they receive when they do.

Wealth Inequality Propagates Financial Risk Taking

Differential participation in stocks, both in taxable and in retirement accounts, tends to produce divergent wealth outcomes, as it gives access to different returns on wealth and ultimately, higher wealth concentration at the top of the distribution where stockholding is more prevalent. For instance, recent research shows that understanding the level and volatility of top wealth shares requires comparing the pattern of historical wealth returns by wealth decile. To date, the wealthy have not only had more to invest, but they also tend to anticipate higher expected returns (Bach, Calvet, Sodini, 2020), and earn higher actual returns on their assets (Fagereng, Guiso, Malacrino, Pistaferri, 2020). This propagates wealth inequality by linking peoples’ current wealth levels to their potential for higher future wealth. Moreover, Haliassos et al. (2023) explored whether those exposed to greater wealth inequality at the start of their careers are more likely to engage in financial risk taking and end up with higher wealth levels later in life; and whether this applies to all or only to some, hence contributing to further widening wealth inequality.

Conceptually, exposure to greater wealth inequality can provide the more able an informative signal about what is possible; to others, it might signal the hopelessness of bridging the gap. Success by the former renders such perceptions self-fulfilling, raising this group’s wealth and influencing subsequent careers. The lack of a positive link between wealth inequality and households’ subsequent wealth outcomes also confirms and reinforces the fact

that people do not respond to inequality by taking more financial risk. As a result, the wealth, and risk-taking gaps between the two groups widen.

To investigate whether such a wealth-propagation mechanism is at work, it would be useful to design an ideal field experiment isolating the role of early exposure to wealth inequality in shaping household risk taking and wealth outcomes later. This would require randomly allocating people at the outset of their economic lives to areas with different local wealth inequality, and then observe their wealth behavior and outcomes in the medium to longer run. While such a thought experiment is not generally feasible, it suggests that the refugee sample from Sweden used in the previous study, could also be of use here. But would using special individuals, such as refugees, be helpful or detrimental to the task?

In our view, the refugee sample is helpful not only for circumventing econometric problems of endogenous selection of areas in which to locate, but also the impact of their neighbors. These refugees were, by definition, at the start of their economic lives in Sweden, so we can focus on refugees starting off destitute (lacking initial wealth heterogeneity). Unlike young Swedes forming impressions about wealth inequality around them throughout their formative years, the refugees were faced with a novel environment and were forming an impression of that environment at the time of entry. Although they came from numerous different countries, researchers knew their country of origin and could control for cultural predispositions, shown by several authors to be relevant for financial behavior. The data also include their education level at entry and later, and the Swedish statistics are carefully recorded to correct for any lack of equivalence of educational levels in origin countries and in Sweden. Finally, the Swedish wealth tax data, LINDA, provide a long list of household characteristics relevant for wealth and portfolio behavior at the time such behavior is observed. Further,

researchers can control for common geographic environment in the form of fixed characteristics of the area of initial allocation and of key time-varying characteristics.

Our study reported that college educated people exposed to greater wealth inequality in their initial locations (measured as the share of total local household wealth owned by the top 10% of the local wealth distribution), were more likely to be successful 10-20 years later in terms of greater wealth, higher position in the wealth distribution of those in their age cohort, and higher ratio of wealth to income. Importantly, these effects of initial wealth inequality at the time of entry were only significant for the college educated located in municipalities with above-median wealth mobility (defined as municipalities where the locals not in the top 10 percent of the wealth distribution had an above-median probability of rising to that top wealth bracket over the 5-year period for which the refugee allocation program lasted). Strikingly, no significant effects of greater initial local wealth inequality were observed for the less educated, even when upward wealth mobility among the less educated locals was considered.

The econometric model we employed controlled for various factors, yet the broad trends were visible even in raw data. Figure 1 focuses on municipalities exhibiting above-median wealth mobility around the time of refugee arrival (1986-1992), splitting each of the two education subsamples by the level of local wealth concentration at the top 10% to which they were exposed upon arrival (above or below median). It computes average wealth of the sampled refugees for each year of wealth observation (1999-2007). The left panel shows the widening wealth gap, even among the college educated, with those exposed to greater local wealth inequality reaching much higher wealth levels than the rest, on average. More striking for social polarization is the comparison to the wealth trajectories of the less educated on the right panel, regardless of initial exposure to local wealth inequality.

Figure 1 here

How did the college-educated end up being wealthier when they were exposed to greater wealth inequality at the start of their economic lives? Our evidence indicates that the college-educated assigned to areas with greater local wealth inequality were systematically more likely to take risks, in the form of stockholding and homeownership, than those educated refugees who were confronted with less wealth inequality. One might expect that those faced with greater local wealth inequality might have been more likely to direct their efforts toward getting better-paying and more secure jobs. Nevertheless, our analysis finds no evidence that being faced with greater wealth inequality generates higher labor income or lower risk of being unemployed, 10-15 years later. Additionally, initial local wealth inequality also did not bear significantly account for refugees' choices regarding whether and how much to invest in human capital following entry.

These findings are consistent with the interpretation that educated people perceive higher wealth inequality as an opportunity in environments that allow mobility into the top of the distribution, and they are more likely to be successful in their risky financial and housing investments, ending up with higher wealth levels. One may wonder whether local wealth inequality simply happens to be proxying for some other aspect of the local environment, which, in turn, is responsible for the observed outcomes. Yet our research includes controls for the fixed characteristics of the municipality across which wealth inequality is measured, as well as time-varying characteristics such as mean wealth, mean income, and income inequality in the municipality. Recall also that significant effects of initial wealth inequality on future portfolio behavior and wealth outcomes are only observed in areas of high wealth mobility. Still, these opportunities apparently fail to mobilize the less educated households and only propel the

college educated to higher levels of risk taking and ultimately, wealth. One might wonder whether the mobility opportunities are tailored to the more educated and this is the reason why the less educated do not respond. Yet the less educated are not significantly influenced by initial wealth inequality even in areas that provide high wealth mobility opportunities for them, as measured by the probability of transition to higher wealth strata.

These findings suggest that the institutional environment is unlikely to be able to explain these differences in mobilization and response to initial wealth inequality. Perceptions of wealth inequality, of the ability to earn portfolio returns and become wealthier, as well as the degree of optimism regarding available asset returns and the nature of social interactions, are likely to be important components of the explanation for the differential pattern of response.

More on Inequality Cues

Recent research on perception, cognition, and developmental and social psychology, focusing mostly on income (rather than wealth) inequality, finds that people receive and process inequality cues from their environment (Phillips et al., 2020; Suss, 2023). These can be social class cues from their peers, information from newspapers, or observation of the built environment. As Sweden is a very transparent society, where local newspapers often report on the wealth of the richest locals, this is more available than in most countries. Yet the least-educated may face greater difficulty accessing these cues, as their peers are less likely to be informed and have direct access to the Swedish language press. It is also possible to show, using cross-municipality regressions, that variation in wealth inequality around the time of entry is well explained by a small number of factors more likely observed by the more educated.

Perhaps more fundamental than perceiving wealth inequality, though, may be the question of whether less educated people are able to make effective use of risky assets to improve their position in the wealth distribution. For one thing, recent research has found that people of low socioeconomic status (i.e., low education and low income) tend to be more pessimistic about asset returns (Kuhnen and Miu, 2017; Das et al. 2019). It has also found that lower education is associated with lower achieved asset returns (e.g., Girshina, 2019). Our research (Haliassos et al. 2020) also suggests that the better educated benefit from financial knowledge of their peers and end up taking more financial risks, while the less educated do not. The differential response may well be traceable to the type of peers with which the two education groups come in contact, and the expectations they foster, as well as to the ability of the two groups to generate good wealth outcomes.

Peers as Advisors

Despite the positive effect of education on outcome discussed above, it does not follow that one could reduce wealth inequality, the differential use of retirement products and premium assets, and social polarization, simply by increasing the number of required school years. Although some research would seem to support an exogenous effect of education on returns, a growing literature finds that educational attainment and returns on wealth tend to be joint outcomes of peoples' innate abilities to process information and to financial acumen (Fagereng, Guiso, Holm, and Pistaferri, 2020; Barth, Papageorge, and Thom, 2020). Instead, it could be promising to explore a multi-pronged, **A.I.D.** approach: **A**lert the less educated about return opportunities and wealth mobility prospects; **I**nform them on how to take advantage of higher-return assets without being destroyed by the risks; and **D**esign simpler financial products with

desirable properties even for less able investors. Empowerment of the less educated in their financial decisions contributes to democratizing finance and is less likely to be distortionary or to raise political opposition than other options, such as wealth taxation: instead of taking resources away from those who succeed, empowerment measures provide means to improve outcomes for those left behind.

While the projects discussed above emphasized interactions with peers, they still faced the fact that no record of the content of such interactions was available. For this reason, the next project we discuss seeks to shed light on what peers tend to recommend to others, and how this compares to what professional advisors would recommend in the absence of reward structures that generate conflict of interest for them. The research was conducted in Germany, a high-income country with a direct stock market participation rate of 15.40% and 20.6% participation in mutual funds.⁴

Prior research has touched on this question. For instance, Rumpf et al. (2024) presented professional advisors with randomly assigned vignettes of investors, to elicit their recommendations on such investors' risky portfolio share for retirement saving. The advisors were told that they would not receive monetary compensation for their recommendations, let alone one linked to the recommended portfolio strategy. Yet they would receive information of interest to them in their practice, conditional on responding to the questions. The objective was to incentivize them to take the exercise seriously, but not to bias their advice. The presentation of vignettes to advisors also overcame econometric issues due to endogenous matching of certain types of clients with financial advisors. Further, it avoided the pressures imposed on financial advisors of having to cater to clients' biases and prior experiences in their effort to win

their business, as would be observed at first meetings of clients with their advisors, even when the clients were in reality ‘mystery shoppers’ employed by researchers.⁵

The study adopted a similar elicitation procedure with lay people ‘from the street,’ subject to the requirements that they had basic knowledge and understanding of financial matters. This was intended to capture the type of advice that people might get by approaching peers who would have something useful to say on the subject. The setup allowed a study of whether and how the elicited advice of professional and lay advisors differed systematically in relation to the investors’ characteristics, to the advisors’ own characteristics, and to having the status of the professional advisor known for given observable advisor characteristics. In all cases, the elicited piece of advice referred to the recommended risky portfolio share in a retirement account. Beyond standard Tobit estimation that allowed for the 0, 1 limit in portfolio shares, the study also employed Bayesian methods allowing for observed and unobserved advisor heterogeneity, to illustrate the (estimated) distribution of advice that different types of investors received from unconflicted financial advisors and from suitably chosen peer groups. The study illustrated what is possible to estimate, by considering heterogeneous advice from one’s age and education peers as implied by ‘homophily’ in forming peer groups; from older peers, proxying the advice from informed parents and other elders; and from younger peers, proxying for the advice obtained by older people from their informed children and their friends.

Controlling for advisor characteristics (but not initially for recipient characteristics), professional advisors tended to recommend a lower allocation to risky assets for retirement than did lay advisors. Older and more risk tolerant advisors, as well as those expecting higher returns over a 10-year horizon, tended to recommend bigger risky exposures. Interestingly, controlling for their characteristics, advisors were positively influenced by their own portfolios in the

advice that they gave to others, encouraging them to tilt portfolios in the same direction of riskiness. In this initial set of regressions, there was still no way to compare the extent of this bias across professionals and lay advisors.⁶

When investor characteristics in the vignettes presented to professional and lay advisors were included in the regressions, both lay and professional advisors tailored their advice to investor characteristics, despite the absence of incentives to cater to specific investor preferences. They tended to recommend higher risk exposures for investors with higher income or wealth or with lower debt. They did moderate their recommendation when the amount to be invested was larger. Advisors also recommended less risk for the more risk averse, for older individuals, and for people with little previous experience in the stock market, while they did not respond systematically to investor educational attainment or marital status. All in all, the study showed that both professionals and lay advisors adjusted their recommendations in the direction implied by portfolio theory, even though their knowledge of such theory cannot be presumed.

When we expand the analysis with interaction terms between the status of professional advisor and the investor characteristics in the vignettes, the moderating effect of professional advice on the risky portfolio share can be fully traced to differences in how professional advisors responded to investor characteristics versus their own, relative to lay advisors. One might expect that professionals would be less responsive to their own characteristics when giving unbiased financial advice. Nonetheless, that study showed that professional advisors were more responsive to their own risk tolerance and their own incomes when recommending risky portfolio shares for others. When it comes to taking characteristics of advisees into account, professional advisors moderated their recommendations with respect to investor's age

and risk tolerance than did lay persons. This contrasts with their willingness to recommend higher risky shares for people with more stock market experience, in contrast to lay advisors who did not adjust their recommendations for those with considerable investor experience.

A further objective of the study was to describe the extent of heterogeneity in advice that a particular potential investor might receive, depending on which advisor the investor asked within a specific advisor group. Advisors, both professional and lay, can differ in ways not fully captured by the advisor characteristics we observe in our survey. To allow Bayesian methods to model unobserved advisor heterogeneity, we ask what posterior distributions characterize the advisors' advice conditional on investor characteristics presented to them in the vignette, to estimate the range and distribution of advice to be potentially received.⁷ This exercise allows researchers to choose investor types that might be of interest to consider, as well as advisor groups that individuals might approach. Of course, this exercise can be conducted for any investor and advisor groupings of interest to researchers or policymakers, if these are defined in terms of the characteristics collected in our survey.

Our study considers a young low earner without college education, a college-educated wealthy retiree, and a wealthy person in the latter half of the working life. The potential advisor types include a professional advisor, a peer with the same education and labor income status, and other advisor options that depend on the investor's age: a more senior family member or peer (in the case of young investors); or a somewhat younger peer who still works, and a considerably younger family member or friend, who could be thought of as an offspring (in the case of older, retired investors). Based on our findings, we show that young low earners with low education are likely to get more conservative portfolio advice if they choose to discuss financial matters with their peers than with older people, but they are likely to get the biggest

risky portfolio share recommendations if they choose to speak to financial advisors. This predicted tendency of financial advisors to outdo the portfolio recommendations of the two other peer groups is reversed for the case of a college-educated wealthy retiree. Professional advisors are likely to be the most conservative towards a wealthy retiree, but the retiree is also likely to be getting more conservative advice from high-income young people than from the own age-education peers. Finally, a wealthy person age 50-65 can expect to get more conservative advice from a randomly chosen professional acting without conflict of interest, than from a randomly chosen peer in the same age-education group. Differences in the distribution of advice are predicted to be small, though, for investors that declare high risk tolerance to their advisors.

The pattern of these results is intriguing, especially in the face of existing research on who tends to be using professional financial advice. For instance, Hackethal et al. (2012) considered samples of clients of independent financial advisors and clients of a major bank in Germany. They showed that the people most likely to be matched with a professional financial advisor were older, wealthier, and more experienced. If we combine the lessons from Rumpf et al. (2024), that professionals not facing conflicts of interest were recommending more limited risk exposure to college educated groups above age 50 versus their peers or their children, the implication is that the current pattern of financial advice, if anything, tends to discourage stockholding exposures. This is reinforced by the further finding in our study that professional advisors not subject to conflict of interest encouraged young, lower-educated individuals to include stocks in their financial portfolios to a greater extent than their peers and elders would advise.

This evidence provides a novel argument in support of the case for promoting access to financial advice for the young and inexperienced with financial markets. That is, the sources this group would normally approach, namely family and friends, would be less likely to encourage them to use stocks and take advantage of the wealth creation potential of the equity premium as part of their portfolio for financing retirement. An important challenge in this context would be to persuade the young that talking to professional advisors can be worthwhile, even with small amounts of savings and potential contributions to retirement plans, while also convincing professionals that the long horizons of such investors compensate for the small initial size of their accounts.

Other Sources of Financial Stress

While the three projects highlighted above focus on the role of others as sources of information, inspiration, or advice, a final project we discuss turns to an analysis of the effects on financial behavior arising from background stress imposed by others that individuals cannot ignore or shake. Such stress can arise from a general crisis (e.g., a pandemic, a fiscal or financial crisis, a war), but also from our daily lives (e.g., stressful work environments, problematic relationships, small children, health shocks). The common element of such stressors we built into the experiment is that they are often on our minds, taxing our ability to concentrate and perform other important tasks. To this end, we use an online experiment on a representative sample of 1,881 French respondents. The specific tasks on which the experiment focused was consumption and saving choices in the presence of labor income risk and, in some treatments, additional uncertainty regarding the occurrence and duration of a significant drop in income (such as what might arise from a furlough or an unemployment spell). Specifically, the question

that we ask is: Does cognitive load, either alone or in conjunction with adverse labor market shocks, affect individuals' consumption and saving decisions and how?

There are many *a priori* reasons to think that interference with people's ability to concentrate and devote their full attention to making sound financial decisions could result in overspending and limited wealth accumulation. This is because people tend to take care of their current needs and underplay the future consequences of their actions [CITE?]. Moreover, certain demographic groups are likely to be more prone to such types of behavior, calling for special attention from policymakers.

Nevertheless, our overall findings turned out to be the opposite. That is, we found that, when confronted with cognitive load that interfered with consumption and saving decisions, people tended to become more cautious, leading them to underconsume and save more than they would have in the absence of that cognitive load, taking into account all their characteristics. Moreover, this pattern was not confined to a few demographic groups but was generally present, with only minor exceptions.

This study involved a 20-period consumption-saving problem in the presence of labor income risk. While the control (group 1) had to make 19 choices (as the 20th was the residual of consuming all remaining wealth), three treatment groups were defined. Group 2 faced in addition a probability that a 30% reduction in period income would be imposed with unknown duration. This can be thought of as furlough or unemployment risk. As expected, members of this group moderated their consumption and raised their accumulated wealth as given this increase in background income risk. The remaining two groups faced a cognitive load when trying to make their consumption/saving decisions: group 3 faced the same labor income process as the control group 1, combined with a cognitive load; and group 4 faced the labor

income process of group 2 incorporating a furlough or unemployment probability with unknown duration.

The cognitive load was incorporated in the form of randomly appearing numbers to which the subjects had to respond within a very limited time, by pressing or not pressing the space bar. To make sure that subjects could not ignore this annoying task, they were told that they would be rewarded for correct answers and that their reward would be multiplicative across the cognitive load and the main (consumption/saving) task. Performance in the main task was also incentivized by instructing participants to keep as close as possible to optimal consumption behavior implied by an expected utility maximization model, reset every period to consider the realized level of available savings, and by continually providing feedback on how they were performing relative to a rational optimizer. Their payoff in the consumption task was linked to the maximized expected utility resulting from their consumption decision in that model period, assuming optimal behavior in the future. In other words, participants were told that they were rewarded based on the best that could be achieved from that model period on, because of their consumption/saving choice in that period. This performance measure was illustrated diagrammatically before they were to make their next choice. To maintain undiminished attention throughout the experiment, subjects were told that, at the end of the task, two model periods would be drawn at random, and they would be rewarded based on their performance in those two periods.

Reassuringly, average behavior of each group tended to get closer to optimal behavior predicted by the respective labor income model as time elapsed and subjects became more familiar with the tasks at hand. Nevertheless, imposing the cognitive load raised average group consumption choices away from optimal behavior, regardless of whether

furlough/unemployment risk was present or not. The ability of the respective rational model to describe heterogeneity in behavior within each group, as measured by the mean squared deviation of actual from optimal consumption, did not monotonically improve as the experiment progressed, but the imposition of cognitive load shifted penalty functions upward. Whether we use the ‘macro’ (group average) measure of model proximity or the ‘micro’ (heterogenous behavior) one, the imposition of cognitive load created substantial additional departures of the subjects from optimal behavior implied by the respective labor income model.

Cross sectional econometric analysis sheds light on the effects of each treatment on average consumption and financial assets chosen by each subject over the life of the model, as well as on the ability of subjects to approximate the behavior implied by the model. Controlling for a range of subject characteristics elicited in the survey, in addition to the experiment, the systematic effect of facing cognitive load was to lower the average consumption choices and to raise the average level of financial assets held by the subjects. This means that responses of subjects facing cognitive load were significantly farther from the rational model predictions, compared to its absence, the bulk of the effect came through a suboptimal level of consumption given the amount of available financial assets (a wrong ‘policy rule for consumption’ in the dynamic programming jargon), rather than from subjects responding optimally to the level of financial assets resulting from previous decisions (a ‘suboptimal evolution of the endogenous state’).

The effects of cognitive load on those not facing increased furlough or unemployment risk were found to be independent of observable subject characteristics. In the example of the Coingvid crisis, this could refer to online workers not facing the prospect of furlough. However, the combination of furlough risk and cognitive load, facing physical presence workers during

the crisis, was potentially a more challenging task for subjects. The study found indeed that college-educated subjects performed systematically better, in that they reduced their consumption less and exhibited smaller deviations from the model than those without college education. Yet it was difficult to detect other subject characteristics that played a role in moderating effects or departures from the model. In fact, one of the few factors that mattered, being patient, resulted in even greater drops in consumption in response to the cognitive load, and in greater deviations from optimal behavior.

This study therefore suggests that people become more cautious in their spending and accumulation of assets during difficult periods occupying their minds, lowering the propensity to spend rather than creating significant problems of household liquidity or insolvency. Providing advice and support on how to relate consumption spending to their resources best could be part of an overall strategy to get the economy going again, without using payments that put pressure on fiscal space. Moreover, one could plan wide-reaching campaigns rather than support targeted to specific disadvantaged groups, as this type of cognitive burden problem seems to apply regardless of demographic and other household characteristics usually considered. Based on the findings of the project, we would also expect our estimated consumption functions to capture behavior less well in the depths of a crisis, as people become more cautious with their resources. On the other hand, when the problems creating cognitive load are individual-specific (such as stressful life events or situations), individuals are likely to require coaching and advice, not in the direction of avoiding overspending but in the direction of adopting a reasonable consumption plan given available resources and avoiding becoming over-cautious.

Conclusions

This chapter has reported new findings from several recent research projects, some of which are still underway. Our key takeaways can be summarized as follows. First, interacting with a larger proportion of neighbors with college-level economics or business education tends to promote retirement saving. Yet people's tendency to associate with others like them means that individuals from low socioeconomic groups are unlikely to be interacting with knowledgeable others, and unable to interpret the information they receive when they do. Second, college-educated people exposed to greater local wealth inequality and wealth mobility at the start of their economic lives tend to take more asset risks later in life and thus achieve greater wealth, leaving the less educated behind. Nevertheless, general education is likely to be insufficient to solve such problems, while empowerment programs of the less educated to improve their financial behavior seem to be needed. Third, the current pattern of access to financial advice, under which the young and less experienced are also less likely to be meeting financial advisors, tends to discourage stock market participation and inclusion of stocks in retirement portfolios, both because peers of the young tend to be more conservative in their recommendations to them, and because financial advisors are more conservative towards their older and wealthier clients. Finally, background stressors such as crises and wars, as well as personal problems, occupy people's minds as they make saving decisions. In an incentivized online experiment, background stress consistently made people more cautious in their financial behavior, in the sense that they reduced their consumption and ended up with larger amounts of financial balances. As there was not much evidence of heterogeneity across individual characteristics, it appears that aggregate crises do not require targeted policies to respond to

these effects. Yet there is reason to support individuals who face person-specific circumstances that prevent them from concentrating on important financial choices for their financial future.

Acknowledgements

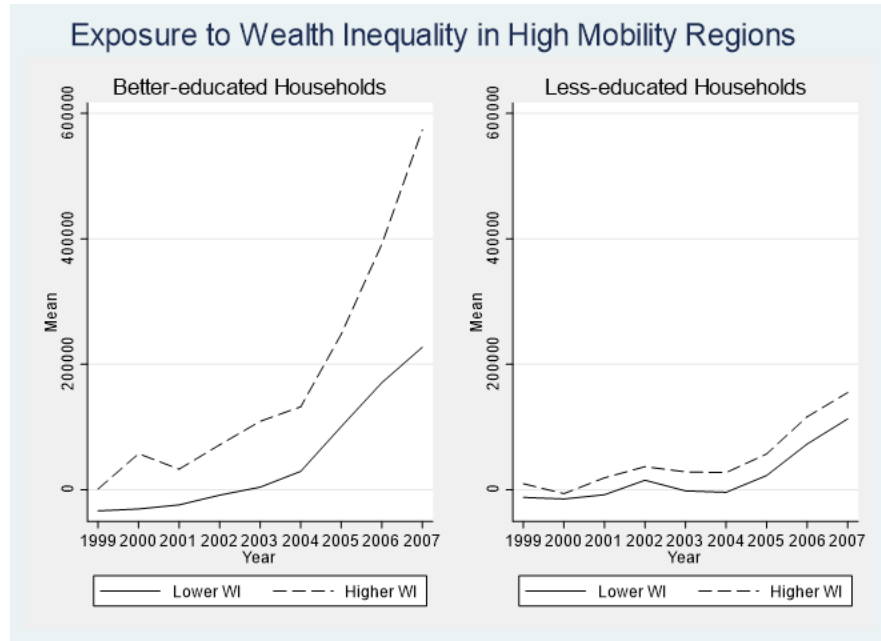
The author thanks his coauthors on projects described including Tiziana Assenza, Alberto Cardaci, Thomas Jansson, Yigitcan Karabulut, Tetyana Kosyakova, Thomas Otter, and Matthias Rumpf. He is also grateful to the funding institutions supporting the work including the European Savings Institute (OEE), the French National Research Agency (ANR) under the ‘Investissements d’Avenir’ program grant ANR-17-EURE-0010, and the German Research Foundation (DFG) [HA 5167/2-1]. He also thanks the colleagues and students, too numerous to list here, who contributed comments or research assistance on the various drafts of working papers associated with the projects. Finally, he is am grateful to Macquarie University in Sydney, which hosted my visit during which part of this paper was written.

References

- Arrondel, L., H. Calvo-Pardo, C. Giannitsarou, and M. Haliassos (2022). 'Informative Social Interactions', *Journal of Economic Behavior & Organization*, 203: 246-263.
- Assenza, T. A. Cardaci, and M. Haliassos (2023). 'Consumption and Account Balances in Crises: Have we Neglected Cognitive Load?', IMFS Working Paper 197.
- Bach, L., L.E. Calvet, and P. Sodini (2020). 'Rich Pickings? Risk, Return, and Skill in Household Wealth.' *American Economic Review* 110(9): 2703–2747.
- Barth, D., N.W. Papageorge, and K. Thom (2020). 'Genetic Endowments and Wealth Inequality.' *Journal of Political Economy*, 128(4): 1474–1522.
- Das, S., C.M. Kuhnen, and S. Nagel (2019). 'Socioeconomic Status and Macroeconomic Expectations.' *Review of Financial Studies*, 33(1): 395–432.
- Diermeier, M., H. Goecke, J. Niehues, and T. Thomas (2017). 'Impact of Inequality-Related Media Coverage on the Concerns of the Citizens.' DICE Discussion Papers 258, Heinrich Heine University Düsseldorf, Düsseldorf Institute for Competition Economics (DICE).
- Fagereng, A., L. Guiso, M. Holm, and L. Pistaferri (2020). 'K-returns to Education.' Working paper.
- Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). 'Heterogeneity and Persistence in Returns to Wealth.' *Econometrica* 88(1): 115–170.
- Girshina, A. (2019). 'Wealth, Savings, and Returns over the Life Cycle: the Role of Education.' Technical Report.
- Gomes, F., M. Haliassos, and T. Ramadorai (2021). 'Household Finance.' *Journal of Economic Literature*, 59(3): 919-1000.

- Hackethal, A., M. Haliassos, and T. Jappelli (2012). 'Financial Advisors: A Case of Babysitters?' *Journal of Banking and Finance*, 36(2): 509-524.
- Haliassos, M., T. Jansson, and Y. Karabulut (2020). 'Financial Literacy Externalities.' *Review of Financial Studies*, 33(2): 950-989.
- Haliassos, M., T. Jansson, and Y. Karabulut (2023). 'Wealth Inequality: Opportunity or Unfairness?' CEPR Discussion Paper 17237.
- Kuhnen, C.M. and A.C. Miu (2017). 'Socioeconomic Status and Learning from Financial Information.' *Journal of Financial Economics* 124(2): 349–372.
- Kraus, M.W., J.W. Park, and J.J.X. Tan (2017). 'Signs of Social Class: The Experience of Economic Inequality in Everyday Life.' *Perspectives on Psychological Science* 12(3): 422–435.
- Manski, C., 1993. 'Identification of Endogenous Social Effects: The Reflection Problem.' *The Review of Economic Studies* 60(3): 531-542.
- Mullainathan, S., M. Noeth, and A. Schoar (2012). 'The Market for Financial Advice: An Audit Study.' NBER Working Paper 17929.
- Phillips, L.T, S.J. Tepper, D. Goya-Tocchetto, S. Davidai, N. Ord- abayeva, M.U. Mirza, B. Szaszi, M.V. Day, O.P. Hauser, and J. Jachimowicz (2020). 'Inequality in People's Minds.' Working paper, Psyarxiv.
- Rumpf, M., M. Haliassos, T. Tesyakova, and T. Otter (2024). 'Financial Advisors versus Lay People: A Horserace' (Working Title).
- Suss, J. H. (2023). 'Measuring Local, Salient Economic Inequality in the UK.' *Environment and Planning A: Economy and Space*, 55(7): 1714-1737.

Figure 1



Source: Haliassos et al. (2023).

Endnotes

¹ Arrondel et al. (2022) found that French households discuss such matters with three to five people on average, while their social circle exceeds 50 on average.

² Unfortunately, Sweden stopped collecting such data following abolition of the wealth tax.

³ In fact, if the refugees were not being influenced by their neighbors directly but by features of their neighborhood, this would be an instance of what Manski (1993) called ‘correlated effects’, and they would not be attributable to interactions with peers.

⁴ This is above average in the Eurozone (10.9% and 12.9%, respectively), but lower than in the US, where direct stockholding rose from 15% to 19% between 2019 and 2022, while combined direct and indirect stock market participation rose from 53% to 58%.

⁵ Evidence of such catering was found by Mullainathan et al. (2012).

⁶ This tendency has been found in existing research on professional advisors, and it is commonly attributed to familiarity bias, that is recommending to others what you yourself know best.

⁷ Specifically, for a given client type, h , with household characteristics x_h , we compute predicted risky asset share recommendations, y_{ah} , for each adviser a of a given peer type. Predicted risky asset share recommendations are computed at each draw of the reaction function of advisor a , defined as the vector of coefficients β_a . The results that are smaller than 0 and larger than 1 are censored to the respective portfolio share limits. All predicted values of risky share recommendations for a given peer type form the distribution of predicted risky asset share recommendations for client h with client characteristics x_h being discussed in the text.