The dynamics of stock market participation

Sigurd Mølster Galaasen† — Norges Bank
Akash Raja* — LSE

August 15, 2022

Please click here for the most updated version

Abstract

We document five new empirical facts about the dynamics of stock market participation using 26 years of Norwegian administrative data: (1) Short spells in the stock market are common, particularly for individuals of low financial literacy, with 22% of all spells ending within 2 years. (2) As time spent in the stock market increases, the probability of exit falls. (3) Re-entry into the stock market is commonplace with > 30% of exiters re-entering within 4 years, and occurs more frequently for individuals with characteristics associated with high financial literacy. (4) Conditional on occurring, re-entry generally happens very soon after exit, often just 1 year later. (5) The likelihood of re-entry falls with time since exit. We show that the workhorse life-cycle portfolio choice model with participation costs fails to produce any of these dynamics, and subsequently propose a model featuring learning about ability and imperfect memory that can quantitatively explain all five facts.

Keywords: household finance, stock market participation, dynamics, re-entry
JEL Classification: D14, D84, G11, G40, G50

† Bankplassen 2, 0151 Oslo, Norway. E-mail: `Sigurd-Molster.Galaasen@Norges-Bank.no`
* Houghton Street, London, UK, WC2A 2AE. E-mail: `a.raja@lse.ac.uk`

We are grateful to Wouter den Haan, Jonathon Hazell, Ethan Ilzetzki, Benjamin Moll, Cameron Peng, Maarten de Ridder and Ricardo Reis, as well as conference and seminar participants at the CEPR Seventh European Workshop on Household Finance, IAAE 2022 Conference, London School of Economics, MFR 2022 Summer Session, MMF 2022 PhD Conference, Norges Bank, RES 2022 Symposium of Junior Researchers and Warwick 2022 PhD Conference, for their valuable comments and suggestions. Raja acknowledges support from the Economic and Social Research Council.

1 Introduction

What determines whether households participate in the stock market? Contrary to the predictions of modern portfolio theory (Samuelson (1969); Merton (1969, 1971)), the data shows that a large proportion of individuals choose not to invest in equities despite their high average premium relative to riskless assets (Mankiw and Zeldes (1991); Haliassos and Bertaut (1995); Campbell (2006)). While various explanations such as participation costs have been proposed for why the participation rate is not 100% at each *static* point in time, much less is known about the decisions to enter into and exit from the stock market. For example, how long do individuals spend in the stock market and, if they leave, do they ever return? This exit and (re-)entry margin is important because these decisions can affect not only household welfare, but also asset prices and macroeconomic aggregates. Furthermore, understanding these dynamics could help to shed light on the wide range of existing theories for (under)participation because they may have conflicting predictions for individual-level movements in and out of the stock market.

This paper uncovers five novel empirical facts on the dynamics of stock market participation using detailed Norwegian administrative data and discusses their implications for theories of participation. The data requirements for studying these dynamics are significant, which can explain the limited existing work on this topic; however, the Norwegian data is particularly well-suited for this analysis relative to alternative datasets. Due to the wealth tax in Norway, the tax records contain wealth information by broad asset class for each individual in the population from 1993 to 2018. The panel structure allows us to reliably identify individual-level movements in and out of the stock market, and the long time dimension is essential for analysing the length and frequency of spells in the stock market. Information on wealth holdings is directly reported by financial intermediaries, thus alleviating concerns about measurement error that can potentially trouble analysis based on wealth surveys. Furthermore, the tax records can be linked to other administrative datasets, thereby providing additional information on each individual that will be useful for identifying heterogeneity in dynamics.

We document the following new facts on participation dynamics: first, very short spells in the stock market are commonplace with 22% of all spells in the data ending within 2 years. We find a higher prevalence of short spells amongst individuals with characteristics typically associated with lower financial literacy, namely low income and wealth and not having a college degree. Second, we apply the methodology of Alvarez et al. (2021) and estimate a downward sloping and convex hazard function for exit from participation. This indicates negative duration dependence in exit probabilities: the longer you have stayed in the market, the lower is the probability of leaving at that point in time. Third, many exiters (>30%) re-enter within 4 years after exit. Contrary to the findings for short spells, re-entrants typically possess characteristics associated with higher financial literacy, in particular high income and wealth. Fourth, conditional on occurring, re-entry generally takes place soon after exit with 45% of re-entry happening just one year later. Fifth and last, we estimate a hazard function for re-entry following exit and find negative duration dependence here too, which means that the longer you have been out of the stock market, the lower is the probability of returning at that point in time. The hazard rate falls sharply during the initial years following exit and by 10 years after exit, the probability of re-entering is effectively zero. Taken together, these facts indicate a high degree of turnover between non-participation and participation states with many individuals having short, multiple spells.

We then provide a discussion of these dynamics in light of the existing broad classes of theories for the underparticipation puzzle and show that these explanations would be unable to generate the short-term dynamics observed in the data. To further verify this claim, we take the workhorse life-cycle portfolio choice model of Cocco et al. (2005) augmented with fixed participation costs and calibrate this model to the Norwegian economy. Model simulations show that the model fails to produce short-term (re-)entry and exits. No spell lasted less than 30 years and there is virtually zero re-entry as a result of most people having just one single long spell throughout their working life.

To rationalise our empirical facts, we augment a Merton model of portfolio choice in three dimensions that reflect human behaviours established in existing work: first, individuals have ex-ante heterogeneous ability. This ingredient is motivated by empirical evidence in Bach et al. (2020) and Fagereng et al. (2020), who find that some individuals are able to generate higher returns on average than others. Second, individuals learn about their ability upon observing their realised returns. Seru et al. (2010) and Linnainmaa (2011) find supportive evidence of such behaviour.[1] Third, following the memory literature which documents that people exhibit an imperfect recollection of past events, individuals in the model recall their experienced return with some noise.

We show that the calibrated model can explain our five empirical facts and performs well quantitatively. Short spells occur (Fact 1) as a result of experiencing poor initial returns, which through the learning process lead individuals to believe that they are low ability investors. Learning about ability also generates a downward-sloping hazard function for exit from participation (Fact 2) because individuals who have been participating for many years should be reasonably confident about their ability and hence require an extremely low return to drive them out of the market, which is very unlikely. In contrast, new entrants have no information and so they are less confident about their own ability. It therefore takes a smaller adverse return to drive them back out of the market. The addition of noisy memory drives re-entry. Re-entry can occur (Fact 3) if individuals receive a positive recollection of their past experiences as this can make individuals more confident and drive them back into the stock market. However, individuals who did terribly in their prior spell should remain non-participants because even with some imperfect recollection, they will still conclude that they are bad at investing. Consequently, most re-entry should occur soon after exit (Fact 4) as it does not take long for individuals with moderately poor returns to receive a positive recollection. However, we should see little re-entry after many years of non-participation because those who remain are likely to have performed poorly in their past spell and are therefore very unlikely to re-enter. This selection also leads to a downward-sloping re-entry hazard function (Fact 5).

Related literature: this paper relates to various strands of the literature: first, it fits into a fairly scarce literature that seeks to understand entry and exit decisions. Bonaparte et al. (2021) use the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF) datasets from the US and show that a large proportion of households enter or exit from the stock market on a biennial basis. Using PSID data, they show that on average, 7.3% (8.7%) of time t households enter into (exit from) non-retirement investment accounts in year $t+2$. The focus of our paper differs in various ways. Bonaparte et al. (2021) show

[1] This notion of learning from experience is in line with a broad literature on experience effects, e.g. Malmendier and Tate (2005); Greenwood and Nagel (2009); Malmendier and Nagel (2011, 2015).

that there is a high degree of turnover; however, they do not study whether this exit is driven by new investors or individuals that have been participating for a long time. Their focus is instead on linking exit and entry decisions to income risk, whereas our focus is on how the likelihood of exit can be linked to time since entry. The overall high turnover documented by Bonaparte et al. (2021) is consistent with our findings - we go further and show that this turnover is driven by individuals having short spells. We also put more emphasis on re-entry decisions. Hurst et al. (1998) and Vissing-Jørgensen (2002) use PSID data and also document a high degree of turnover in risky financial markets.[2] Interestingly, Vissing-Jørgensen (2002) finds a non-negligible share of households who enter the stock market sometime between 1984 and 1989, but then leave again at some point between 1989 and 1994. This finding links closely to our first empirical fact of short spells, although the wide time gap of five years between PSID data waves means that in principle, some of these households could have had a spell as long as 10 years by entering in 1984 and leaving in 1994. We are able to narrow the time intervals using annual data and find that a non-negligible proportion of spells end within just 2 years. Furthermore, we study the characteristics linked to short spelling and also consider re-entry.[3]

Second, we draw from work on learning and experience effects. There is a broad literature that documents individuals responding to past experiences across a range of settings, both financial (e.g. Kaustia and Knüpfer (2008); Chiang et al. (2011); Malmendier and Nagel (2011, 2015); Knüpfer et al. (2017)) and non-financial (e.g. Alesina and Fuchs-Schündeln (2007); Oreopoulos et al. (2012)). Seru et al. (2010) use Finnish transaction data to analyse whether individuals learn from their trading experiences. The authors distinguish between two forms of learning, namely learning from experience and learning about ex-ante ability, and find that most learning is of the second type whereby low ability investors stop trading after performing poorly. Linnainmaa (2011) also uses Finnish data and finds that investors increase their trade sizes following successful trades, but exit following poor realisations.[4] We apply these ideas in our model and show that having participants learn about ex-ante

[2] Hurst et al. (1998) correlate entry decisions with observable characteristics and find that income, race and education are predictive of the decision to become a stockholder. However, their focus is not on spell lengths, nor re-entry.

[3] Other papers have studied dynamics across other dimensions such as the life-cycle (Poterba and Samwick (1997); Ameriks and Zeldes (2004); Fagereng et al. (2017a)), house purchases (Brandsaas (2021)) and portfolio characteristics (Calvet et al. (2009a)).

[4] See also Nicolosi et al. (2009).

ability is able to generate the short spell behaviour and downward-sloping hazard function of exit we uncover in the data.[5]

Third, we relate to papers on belief and returns heterogeneity. The basic Merton (1969) model shows that differences in the expected risk premium can generate different portfolio choices. Dominitz and Manski (2011) find evidence for heterogeneous beliefs of equity returns in the US. Hurd et al. (2011) and Hudomiet et al. (2011) find, using Dutch and US data respectively, that those with higher returns expectations are more likely to participate in the stock market. We apply these ideas in our model by allowing individuals to have time-varying expected returns with the variation resulting from changes over time in the belief of being a low ability investor. Our model thus requires that individuals have different innate abilities in the stock market. This is supported by the literature on returns heterogeneity.[6] Fagereng et al. (2020) use Norwegian data and find that that a non-negligible proportion of variability in returns can be explained by persistent heterogeneity. Bach et al. (2020) also find an important role of type dependence in determining returns using Swedish data.[7]

Last, we apply ideas from the literature on (imperfect) memory. Azeredo da Silveira et al. (2020) study optimal memory structure when memory storage is costly, and show that it is optimal for individuals to recall with noise a single summary statistic of their past experience. We use this idea by having investors recall the average of their experienced returns with noise. An outcome of our model is that those investors who return to the stock market are typically those who did slightly poorly in their prior spell(s) rather than those who did terribly. This is in line with papers suggesting that individuals will particularly remember more salient events (e.g. Brocas and Carrillo (2016)).

Outline: The paper is structured as follows: Section 2 describes the Norwegian data, while

[5]While our model applies rational Bayesian learning in line with Seru et al. (2010) and Linnainmaa (2011), there is a literature suggesting that agents may learn differently from standard Bayesian updating (see, amongst others, Charness and Levin (2005); Barberis et al. (2018); Barber et al. (2019); Kuchler and Zafar (2019); Anagol et al. (2021)).

[6]Gabaix et al. (2016) propose "type dependence", whereby individuals have different ex-ante ability in generating returns (e.g. due to education or innate talent), as a possible explanation for the observed positive correlation between wealth and returns documented in various papers (e.g. Bach et al. (2020); Fagereng et al. (2020); Xavier (2021)).

[7]Returns heterogeneity has also been uncovered in the literature on excessive trading, e.g. Barber et al. (2008) show that investors who trade more frequently tend to earn lower returns after fees. Furthermore, the literature on financial literacy shows that individuals with higher cognitive ability perform better in the stock market (Grinblatt et al. (2011)).

Section 3 goes through each of the empirical facts. Section 4 discusses existing theories of participation and other candidate explanations, while Section 5 details our proposed model. Section 6 concludes.

2 Data

We use Norwegian administrative data to conduct our analysis. Unlike most administrative datasets which only have information on income, the Norwegian data also has detailed information on wealth holdings due to the existence of a wealth tax in Norway, thus allowing us to study participation in the stock market. These tax records contain information on assets and liabilities for each Norwegian resident as of December 31st of each year from 1993 to 2018. We are also able to merge the tax records with data containing demographic characteristics.

The Norwegian administrative data is more suitable than other datasets for studying the dynamics of stock market participation for a variety of reasons: first, to study dynamics, we need to be able to follow individuals over time. Our data provides this panel dimension unlike many survey-based datasets, which tend to be repeated cross-sections. Second, while other datasets such as transaction-level data (e.g. brokerage data used in Barber and Odean (2000, 2001)) or non-Norwegian administrative datasets (e.g. Swedish data used in Calvet et al. (2007, 2009a,b)) provide a panel structure, we need a sufficiently long time dimension to be able to study spell lengths and re-entry. These alternative datasets typically do not span a large number of years.[8] Third, a concern with brokerage accounts data is that an individual's exit from the data does not necessarily mean exit from the stock market. For example, if you switch brokers, you would appear as an exiter in the brokerage data, but in reality you are still in the stock market. In addition, re-entry may be difficult to identify if account numbers change between spells. The Norwegian data does not have this concern as the tax data is based on overall holdings across all broker firms and identification is at the individual level. Fourth, further concerns with brokerage data are potential sample selection and non-random attrition, the latter of which is also a significant concern with panel sur-

[8] As the wealth tax in Sweden ended in 2007, the Swedish data used in Calvet et al. (2007, 2009a,b) spans just 8 years (1999-2007). The brokerage data of Barber and Odean (2000, 2001) covers 5 years from 1991-1996.

vey data. Are customers of this particular firm representative of all investors? Do customers leave the sample is a non-random way? Such concerns are much less prevalent in our data as we study the entire Norwegian population and attrition should be only due to death or emigration. Fifth, an issue with using survey-based data is measurement error. Individuals may not perfectly understand the questions or not know their exact wealth holdings perfectly.[9] For our purposes, it would have serious implications if respondents "forget" their exposure to risky financial assets in one survey wave that they then remember in the subsequent wave as this would appear as an exit followed by re-entry rather than one continuous spell. A major advantage of the Norwegian administrative data is that financial institutions directly report information on wealth holdings to the tax authority. Following this direct reporting, residents are sent a pre-filled tax form to approve. If they do not respond, then the tax authority assumes the information is correct and this dictates their tax calculation.[10] As such, it is very difficult to evade taxes in Norway via underreporting of wealth holdings.[11] Last, we are able to link these tax records to other administrative datasets containing information on demographics, employment and house purchases. This allows us to see whether the behaviours we observe are linked to certain characteristics. Such information is not necessarily available in survey or brokerage data.

While the Norwegian data is particularly promising for our research objective, it has its shortcomings. The data gives us asset holdings as of December 31st of each year. As such, we are limited to participation decisions at the annual frequency, although it is worth noting that this is more frequent relative to most panel surveys.[12] This also means we are unable to capture within-year spells, although the presence of within-year spells would strengthen our result that short spells in the stock market exist. In addition, we do not have information on occupational or public pension wealth; however, in Section 4.5.3 we discuss how pensions

[9] Lusardi and Mitchell (2011) use three simple questions on compound interest, inflation and financial diversification to elicit financial literacy, and find that only one-third of respondents could answer all three questions correctly. This illustrates the potential difficulties many survey respondents may have when confronted with finance-related questions.

[10] In 2009, around 60% of tax payers in 2009 did not respond (Fagereng et al. (2017a)).

[11] As noted in Fagereng et al. (2017a), one source of under-reporting could be if individuals hold but fail to disclose foreign investments. While asset holdings through Norwegian financial intermediaries are directly reported, this is not the case for foreign holdings. For Sweden, Calvet et al. (2007) argue that such holdings are likely to be a small portion of overall assets other than for the wealthiest individuals.

[12] For example, wealth information in the PSID was only collected from 1984 and at five-year intervals until 1999, when it was added to the main interview and then collected biennially.

are unlikely to be affecting our results.

2.1 Data construction

Here we give a broad overview of the construction of wealth variables, in particular our measure of overall participation in risky financial markets, from the tax records. We construct measures of wealth by broad asset class and combine them to obtain measures of *financial* and *real* wealth. Financial wealth can be decomposed into the following asset classes: (a) cash and deposits (both domestic and foreign), (b) directly-held listed stocks, (c) directly-held unlisted stocks (typically private equity), (d) stock mutual funds, (e) money market funds, (f) financial wealth held abroad and (g) other financial assets.[13] Real wealth consists of housing and other real assets.[14] We are most interested in the extensive margin of participation and treat an individual to be participating in a given year if any of directly-held listed stock holdings, stock mutual fund holdings or financial wealth held abroad are strictly positive.[15] We provide further details on the construction of these asset classes from the tax records in Appendix A. The only sample selection criterion we impose is only looking at individuals aged 20 or over to ensure that the person is the main asset holder.[16]

2.2 Descriptive statistics

Table 1 gives summary statistics at the individual level for the pooled sample from 1993-2018, which spans over 97 million observations. The first block shows that there is an even split of males and females in the sample. 28% of individuals have a college degree. The second block provides information on income and wealth holdings. The average individual has a total gross wealth holding of \$217,000, though the large standard deviation in asset holdings illustrates the wide heterogeneity in wealth across the population. In particular, the median wealth holding is less than half of the mean holding, indicating a rightward skew

[13] Other financial assets consists of outstanding claims and receivables, shares of capital in housing cooperatives or jointly-owned property, own pension insurance and life insurance, and other wealth.

[14] Other real assets include vehicles (e.g. boats, cars, caravans), holiday homes, fixtures and other business assets, contents and other real estate (e.g. farms, plots).

[15] We include financial wealth held abroad in this definition to be conservative because the nature of such wealth is not observed. However, few people hold wealth abroad ($<2\%$ of observations).

[16] Fagereng et al. (2020), in their study of heterogeneity in the returns to wealth across the wealth distribution, also impose an upper bound on age of 75. However, we do not do this as it can artificially generate right-censored spells.

in the wealth distribution. Non-financial wealth, of which a major component is housing, accounts for a larger share of total wealth than financial wealth with the average individual holding \$64,000 in financial wealth compared to \$153,000 in non-financial wealth. The mean amount of wealth held in public equity, measured as the sum of holdings in stock mutual funds, directly-held stocks and financial wealth abroad, is just under \$8,000. Indeed, the median individual does not hold any public equity, a finding that is indicative of broad aggregate underparticipation in the stock market in Norway.

TABLE 1: Summary statistics

	Mean	Std. dev	P10	Median	P90	P99
Demographics						
Age (in years)	48.48	18.41	25.00	46.00	75.00	90.00
Male	0.50	0.50	0.00	0.00	1.00	1.00
Single	0.37	0.48	0.00	0.00	1.00	1.00
College degree	0.28	0.45	0.00	0.00	1.00	1.00
Income and wealth (2011 \$'000s)						
Gross income	42.52	55.75	0.00	37.02	93.35	191.72
Financial wealth	64.46	1,576.90	0.05	9.98	110.60	666.07
Financial wealth in public equity	7.98	351.36	0.00	0.00	8.08	128.28
Non-financial wealth	153.21	272.87	0.00	67.67	406.08	1,006.48
Gross wealth	217.67	1,651.38	0.29	105.75	495.22	1,469.05
Net wealth	209.40	1,854.63	0.80	99.09	457.76	1,458.58
Participation and wealth shares						
Participates in public equity	0.26	0.44	0.00	0.00	1.00	1.00
Participates in mutual funds	0.21	0.40	0.00	0.00	1.00	1.00
Participates in indiv. stocks	0.10	0.30	0.00	0.00	0.00	1.00
Public eq. share (of gross wealth)	3.03	11.08	0.00	0.00	6.08	65.64
Public eq. share (of fin. wealth)	7.97	19.57	0.00	0.00	30.93	92.65
Observations	97189499					

Note: this table provides summary statistics based on the pooled sample from 1993-2018. The first block gives summary statistics for demographic characteristics. "Single" is a binary variable equal to 1 if the individual is neither married nor cohabiting, and zero otherwise. The second block information on income and wealth measured in 2011 USD (in thousands) based on an exchange rate of \$1=5.9927 NOK at the end of 2011. "Gross income" is income from all sources. "Public equity" is measured as the sum of holdings in stock mutual funds, directly-held stocks and financial wealth abroad. The third block gives summary statistics on stock market (i.e. public equity) participation and the share of wealth invested in public equity.

The third block further verifies this by showing that 26% of observations correspond to participation in stock markets. Most participants invest in mutual funds rather than directly holding stocks. Figure D.1 plots a time series of the stock market participation rate in Norway. The participation rate accelerates during the 1990s for reasons including improved access to financial markets for retail investors, the rise of mutual funds and the growing interest in technology stocks during the dot-com bubble.[17] Interestingly, the participation rate has shown a steady decline from a peak of 33% in 2001 to around 25% by 2018, which on first glance appears to contradict the notion that participation costs govern participation decisions given that technological innovations have made participation in the stock market easier for retail investors. Figure D.3 plots the entry and exit rates over time and shows that the entry rate into the stock market has fallen by half from 5% in 2000 to around 2.5% in 2018, which could contribute to this drop in the overall participation rate over this period.

3 Empirical facts

3.1 Fact 1: short spells are common, particularly among low financial literacy groups

We begin by looking at the distribution of spell lengths in the data. Figure 1 plots a histogram with the distribution of spell lengths based on spells beginning between 1994 and 2015 inclusive.[18] We choose to restrict attention to spells starting by 2015 to ensure that participants have at least three years in which to exit. If, for example, 2017 entrants were also included, they would either have a 1-year completed spell or be right-censored, and so including such entrants could artificially inflate the bars corresponding to a short spell length. The histogram shows a declining relationship between spell length and the proportion of observations. Almost 15% of all spells end in just 1 year and 23% end within 2 years. Indeed, this seems to immediately contradict the notion of participation being driven by on-entry participation costs such as time taken to set up an account. If participation were simply governed by such costs, we would expect to see individuals staying in the market for long periods of time following entry.

We undertake various robustness checks to be able to safely conclude that short spells are commonplace: first, one may be worried that the low proportion of long spells simply

[17] Figure D.2 shows the participation rates separately for stock mutual funds and directly-held stocks. Participation in mutual funds rose by more than fivefold from 1993 to the early 2000s. Participation in directly-held stocks also rose, but by a smaller margin from just over 8% in 1993 to around 12% in 2000.

[18] Left-censored spells are excluded from this figure as a spell length cannot be computed for such spells. These spells are typically those that were already ongoing at the start of the data sample in 1993, though other reasons for left-censoring could be immigration of an existing stockholder into Norway.

FIGURE 1: Distribution of spell lengths

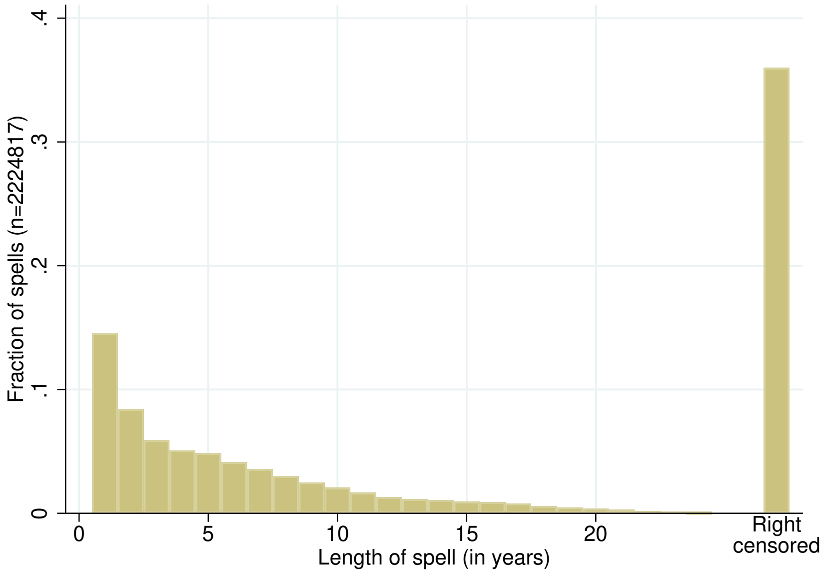

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data. We take all spells beginning at any point from 1994-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells (n=2.2m) belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

reflects the fact that there are fewer ways to have such spell lengths in the data. For example, the only way to have a 24-year spell is to enter in 1994 and to leave in 2018. To address this, Figure D.4 shows the proportion of all spells ending within 1, 2 and 3 years accounting for right-censoring by taking all spells beginning by 2017, 2016 and 2015 respectively. We reach the same conclusion: 14.6% of all participation spells starting by 2017 end within 1 year and 23.1% of spells starting by 2016 end within 2 years. Second, the analysis is done at the individual level and so perhaps the short spells simply reflects a transfer of ownership of financial assets between spouses. We therefore produce the same figure as Figure 1 but at the household level in Figure D.5, and obtain similar results.[19] Third, one may be worried that short spells reflect people receiving a gift or inheritance containing public equity that they instantly liquidate. While we cannot directly see the specific items received, we apply three robustness checks to try to address this concern. We use tax records information to exclude individuals receiving a gift or inheritance above 10,000 NOK (≈ $1670 using 2011 USD)

[19]A household is said to be participating in the stock market in year t if at least one spouse has some assets held in public equity.

in the year of or before entry. We also exclude entrants for whom a parent or grandparent died in the year of or before entry, and in the third check, we exclude entrants for whom a parent or grandparent held risky financial assets in the year of or before entry. As shown in Figure D.6, these robustness checks generate very similar histograms to our baseline figure. Fourth, one might be worried that short spells are driven by individuals holding stocks in the company they work for, which they perhaps sell upon changing jobs. We make use of the Shareholder Registry and demographic information about place of work to identify entrants who hold stocks in their company of work. As the Shareholder Registry data is only available from 2004, this analysis is based on a narrower sample from 2004-2015. Figure D.7 shows that this subsample does not drive the short spells result. Last, we show that quick exits are not driven by entrants who invest small sums of money. Figures D.8a and D.8b show that the findings are robust to restricting attention to individuals who invest at least \$100 and \$1000 at the point of entry respectively.

The next step is to understand whether there is heterogeneity in the prevalence of short spells based on observable characteristics. To do this, we estimate the following linear probability model:

$$\Pr(\text{spell ends within 2 years}) = \alpha_i + \delta_t + \beta' X_{it} + \epsilon_{it} \tag{1}$$

where δ_t denotes entry year fixed effects, and X_{it} is a vector of observable characteristics measured at the point of entry such as age and wealth. Given that we observe individuals with multiple spells, we are able to include individual fixed effects α_i to absorb (unobserved) time-invariant characteristics.

Table 2 shows the result from this estimation in specifications with and without individual fixed effects. Having a partner reduces your probability of a short spell in both specifications. In addition, entrants who enter into directly-held stocks (rather than mutual funds) are 9.1pps more likely to have a short spell. Characteristics that are typically associated with lower financial literacy are also linked to a higher prevalence of short spells.[20] Not having a college degree is associated with a 2.3pp lower probability of exiting within 2 years following entry. Figures 2a and 2b plot the coefficients on the income and wealth decile fixed effects

[20] Lusardi and Mitchell (2011) give evidence of a positive correlation between educational attainment and financial literacy. Behrman et al. (2012) find this too and also show a positive correlation between wealth and financial literacy.

respectively.[21] For income, we can see a monotonic negative relationship between income and the probability of a short spell with those in the bottom income decile having a 2.5pp higher probability of a short spell relative to the median income group. For wealth, the impact of low wealth is even more striking. Entrants belonging to the bottom wealth decile are 10pps more likely to exit within 2 years relative to the median group. Taken together, it appears that short spells are more prevalent for individuals with characteristics linked to lower financial literacy. By age, we see that short spells are more likely for the youngest and oldest age groups (Figure 2c).

It is important to note, however, that this does not mean that short spells are exclusive to these subgroups. Indeed, Figure D.9 shows the distribution of spell lengths by income, wealth, education, gender and asset class. We can see that while the higher prevalence of short spells for certain subgroups shown in the regressions still exists in these figures, there is still a non-negligible proportion of short spells amongst the other subgroups too. As such, short spells are widespread and not purely concentrated amongst a particular subpopulation.

3.2 Fact 2: downward sloping hazard function for exit from participation

Are investors more likely to exit the stock market in the initial periods following entry or after staying in the market for a prolonged period? To answer this question, we estimate the hazard function for exit from participation. The hazard function $h(d)$ gives the probability of exiting the market d years after entry conditional on not exiting until then. A standard challenge with hazard function estimation is separating true duration dependence from (unobserved) heterogeneity. As noted in Lancaster (1979) and Kiefer (1988), estimating hazard functions based on pooled samples with heterogeneous individuals can lead to a downward bias in the slope of the hazard function. If individuals have different underlying propensities to "survive", individuals who are less likely to survive will exit the sample earlier than others. This dynamic selection would bias the hazard function downwards.[22]

[21] In Figure 2a, there is no 2$^{\text{nd}}$ decile for income. This is because > 20% of observations have zero income, and these are all grouped in the first decile. As such, the first decile can be thought of as a zero income group. This will also be the case in later plots of coefficients for income deciles.

[22] Unobserved heterogeneity has been a worry for estimating hazard functions in other settings such as unemployment duration (see e.g. Kroft et al. (2016); Mueller et al. (2021)) and price spell duration (see Nakamura and Steinsson (2008)). For unemployment, the concern is that less employable workers select into long-term

TABLE 2: Determinants of short spells (≤2 years)

Male	0.047***	
	(0.001)	
College degree	-0.008***	-0.023***
	(0.001)	(0.004)
Homeowner	0.001	-0.009**
	(0.001)	(0.003)
Unemployed	0.017***	0.000
	(0.001)	(0.003)
Single	0.023***	0.016***
	(0.001)	(0.002)
Directly-held stocks	0.083***	0.091***
	(0.001)	(0.002)
Sample mean	0.23	0.36
Individual FE	No	Yes
Entry year FE	Yes	Yes
Age group FE	Yes	Yes
Income decile FE	Yes	Yes
Wealth decile FE	Yes	Yes
Observations	2242427	866406
R-squared	0.04	0.47

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. This table shows the estimation of Equation 1. The first column excludes individual fixed effects, while the second column includes them. The dependent variable is a binary variable equal to 1 if the spell ends within 2 years, and zero otherwise. *Homeowner* is a binary variable equal to 1 if the participant owns their own property (either self-owned or ownership through housing cooperatives), and zero otherwise. *Single* is a binary variable equal to 1 if the participant is neither married nor cohabiting, and zero otherwise. *Unemployed* is a binary variable equal to 1 if the participant receives unemployment benefits at the point of entry, and zero otherwise. *Directly-held stocks* is a binary variable equal to 1 if the participant buys directly-held stocks at the point of entry, and zero otherwise. Entry year fixed effects are included. Age fixed effects by broad age group (20-29, 30-39, 40-49, 50-59, 60-69 and 70+), as well as income and wealth decile fixed effects are included. Observables are measured at the point of entry. Standard errors are clustered at the individual level. The regression uses data on entrants from 1994-2016.

To address this concern, we apply the linear GMM estimator of Alvarez et al. (2021) to estimate a discrete time proportional hazard model of duration allowing for unobserved heterogeneity. Under the proportional hazards model, it is assumed that the hazard rate is given by $h_i(d) = \theta_i b_d$. θ_i is the time-invariant frailty parameter specific to individual i and captures individual heterogeneity in hazard rates. b_d is the baseline hazard at duration d and is assumed to be common across individuals. The objective is to obtain an estimate of b_d as this reflects true duration dependence rather than unobserved heterogeneity. The Alvarez et al. (2021) estimator gives a consistent estimator of the baseline hazard when we have panel data on a large number of individuals and we observe at least two spells for some

unemployment, while for prices, products with fairly inflexible prices will select into the group of products with long spell durations.

FIGURE 2: Impact of income, wealth and age on the probability of a short spell

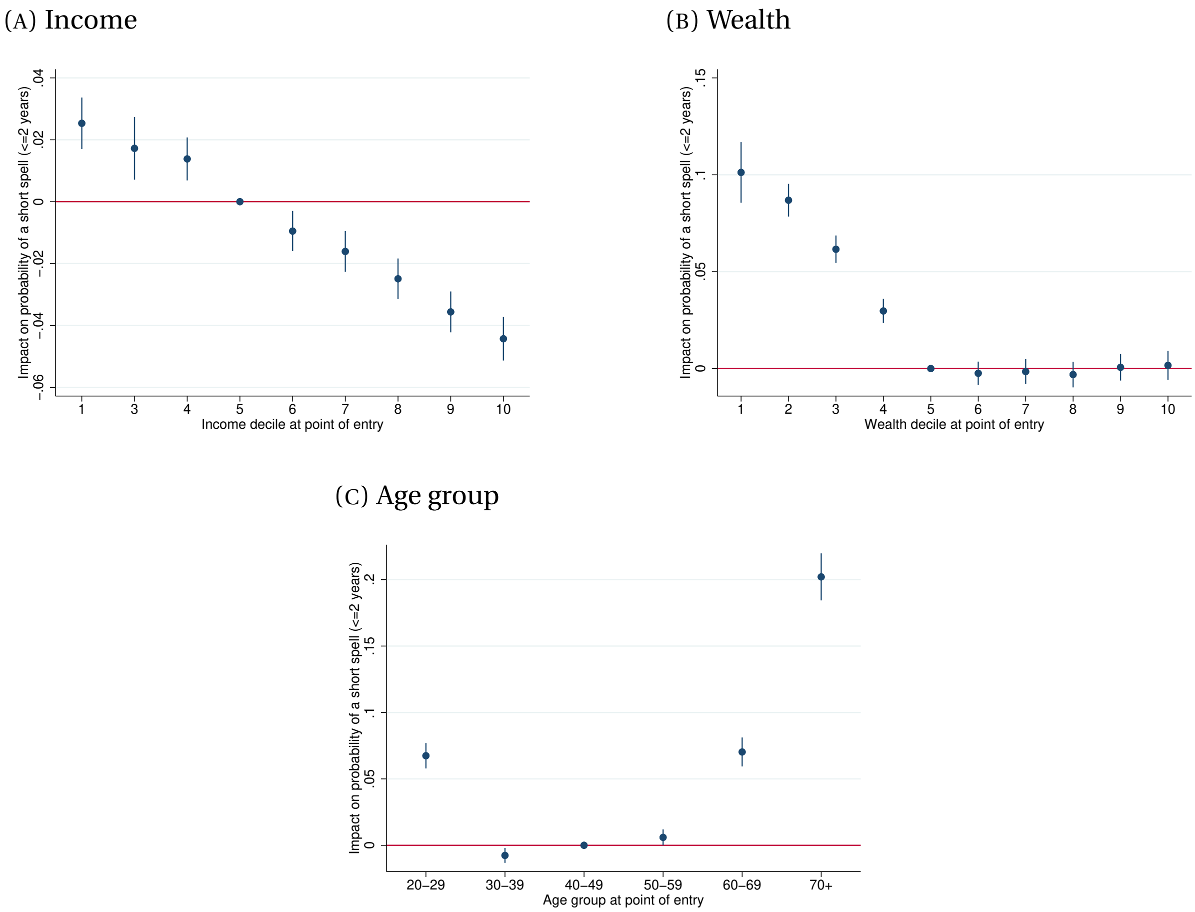

Note: this figure plots the coefficient estimates for the fixed effects on income and wealth deciles following estimation of Equation 1. Individual fixed effects are included in this specification. Variables are measured at the point of entry, and deciles are based on the full Norwegian population aged 20 and above in that year. Panel (A) shows the average marginal effects of income and panel (B) shows the impact of wealth. The effects are estimated relative to the median (5^{th} decile). Panel (C) gives the average marginal effects of age. 95% confidence intervals are shown. The red line represents a null relative effect.

individuals. As such, the estimator relies on some individuals having multiple spells. We will show in Section 3.3 that this is the case in our setting for a non-negligible proportion of participants. There are various advantages of this approach: first, while Honoré (1993) provides continuous time identification results for duration models with multiple spells, the moment conditions used in the GMM estimator are based on discrete time identification results and so this approach is well suited to the discrete time nature of our dataset. Second, some approaches rely on specification of a frailty distribution. For example, Nakamura and

Steinsson (2008) apply the empirical model of Meyer (1990) model in their analysis of price spell duration and assume that the frailty parameter follows a gamma distribution for their baseline specification. Heckman and Singer (1984) note that misspecification of the frailty distribution can bias the hazard function. Instead, the approach of Alvarez et al. (2021) impose no restrictions on the frailty distribution. Third, the GMM estimator is consistent when the number of individuals is large, but it allows for a short time dimension. The latter is important in our setting given that we rely on annual data covering 26 years. Details on the moment conditions and estimation procedure are given in Section C.1.

Figure 3 plots the estimated baseline hazard function. The hazard function is monotonically declining in duration, indicating negative duration dependence, i.e. the longer you have been participating for in the stock market, the lower is the probability of completely exiting at that point in time. As described in Section C.1, we are able to recover the baseline hazard up to a multiplicative constant. As the hazard rate for $d = 1$ is normalised to 1, the hazard rate values give the hazard rate at duration d relative to a duration of 1 year. A striking feature of the hazard function is the steepness of the slope in the initial years following entry. The hazard rates at $d = 2$ and $d = 3$ are 60% and 40% that of $d = 1$ respectively. By $d = 12$, the hazard rate is very close to zero, suggesting that if you have remained in the market for a prolonged period of time, the likelihood of you completely exiting the market is very small. Combined with Fact 1, this indicates strong dynamics in the initial years following entry with a large degree of exit coming from individuals who recently entered the market.

3.3 Fact 3: re-entry does occur, particularly for individuals with high financial literacy

We now turn to understanding whether multiple spells occur - do exiters re-enter following exit and if so, what characteristics correlate with the likelihood of re-entering? Figure 4 plots the distribution of the number of spells an individual experiences. In Figure 4a, we consider the full Norwegian population, while in Figure 4b we look at the distribution conditional on having at least one spell. In both cases, we restrict attention to individuals who appear in the data for at least 15 years as those who appear for fewer years are likely to have either zero or

FIGURE 3: Baseline hazard function for exit from participation

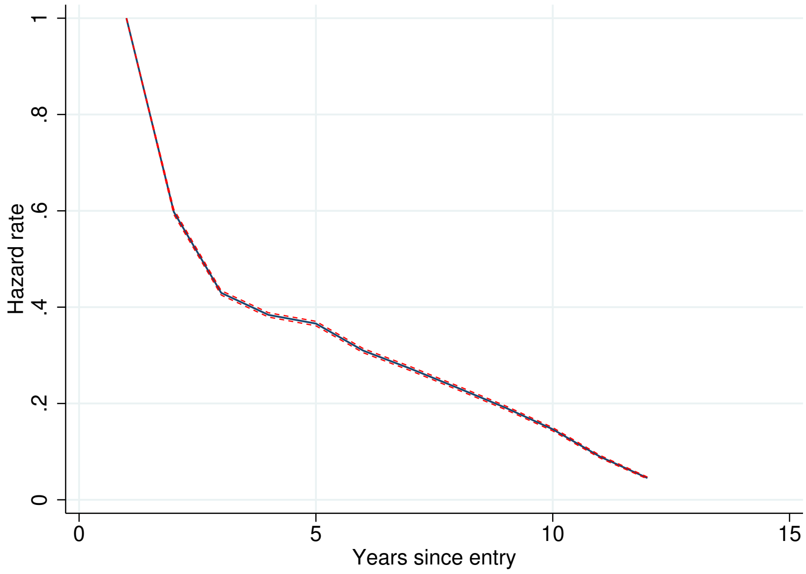

Note: this figure plots the estimated baseline hazard for exit from participation following the methodology of Alvarez et al. (2021) described in Section C.1. The dotted red lines denote 95% confidence intervals. The hazard rate at duration $d = 1$ is normalised to 1.

one spell, which would bias the distribution to the left. From Figure 4a, we see that over half of individuals have at least one spell in the stock market. Given average participation rates of around 25-30% in recent years (see Figure D.1), the fact that over 55% of individuals have at least one spell indicates that there is movement in and out of the stock market, i.e. individuals do transition between non-participation and participation states. Figure 4b shows that amongst the set of individuals who have at least one spell, just under 30% have multiple spells. From this, we conclude that re-entry does occur for a non-negligible proportion of participants.

We now ask which characteristics are associated with re-entry. To study this, we run the following linear probability model:

$$\text{Pr(re-enter within 4 years)} = \alpha_i + \delta_t + \beta' X_{it} + \epsilon_{it} \tag{2}$$

where δ_t denotes exit year fixed effects and X_{it} denote observable characteristics. As some individuals have multiple spells out of the stock market, we are able to include individual

FIGURE 4: Number of spells

(A) All individuals (B) Participants

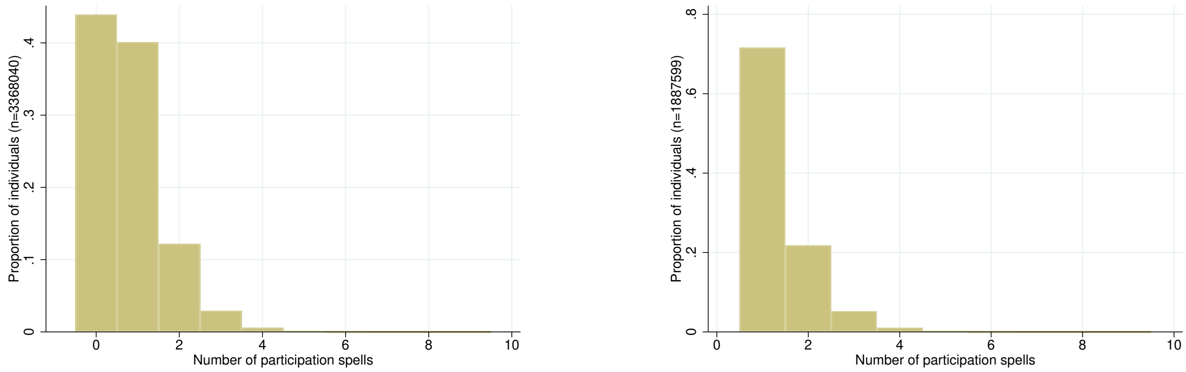

Note: this figure plots the distribution of the number of spells. In panel (A), we use the full Norwegian population, while in panel (B), we restrict attention to those individuals who have at least one spell. In both cases, the individual must appear in the sample for at least 15 years.

fixed effects α_i to capture unobserved time-invariant heterogeneity. We use a fixed window of 4 years to re-enter because those who exit early in the sample have more years remaining in which they could re-enter. A fixed window thus allows all exiters to have the same amount of time in which to re-enter. Furthermore, to preview the findings in Fact 4, we show that most re-entry occurs soon after exit, often just 1 year after, and so a 4-year window should capture a large proportion of re-entry. To ensure that all exiters have at least 4 years in which to re-enter, we restrict attention to those who exit by 2014, four years before our dataset ends.

Table 3 gives the regression results in specifications with and without individual fixed effects. Being a homeowner and single lower the re-entry probability by 2.2pps and 5.6pps respectively. Figure 5 plots the estimated effects of income, wealth and age. The top income and wealth deciles have a higher probability of re-entry. The very top income decile is 4pps more likely to re-enter relative to the median income group (Figure 5a), and the highest wealth decile group is about 8pps more likely to re-enter relative to the median wealth group (Figure 5b). These results together suggest that re-entry is positively associated with characteristics linked to higher financial literacy. Re-entry is less likely for the youngest and oldest age groups (Figure 5c), the latter of which is in line with the finding in Fagereng et al.

(2017a) that permanent exit rises sharply after retirement.

TABLE 3: Determinants of re-entry

	Re-entry in 4y	
Male	0.037***	
	(0.001)	
College degree	0.030***	0.011
	(0.001)	(0.007)
Homeowner	-0.063***	-0.022***
	(0.001)	(0.004)
Unemployed	-0.005**	0.001
	(0.002)	(0.004)
Single	-0.028***	-0.056***
	(0.001)	(0.003)
Sample mean	0.35	0.59
Individual FE	No	Yes
Exit year FE	Yes	Yes
Age group FE	Yes	Yes
Income decile FE	Yes	Yes
Wealth decile FE	Yes	Yes
Observations	1436019	518995
R-squared	0.14	0.54

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. This table shows the estimation of the linear probability model in Equation 2. The first column excludes individual fixed effects, while the second column includes them. The dependent variable is a binary variable equal to 1 if the exiter re-enters within 4 years following exit, and zero otherwise. *Homeowner* is a binary variable equal to 1 if the participant owns their own property (either self-owned or ownership through housing cooperatives), and zero otherwise. *Single* is a binary variable equal to 1 if the participant is neither married nor cohabiting, and zero otherwise. *Unemployed* is a binary variable equal to 1 if the participant receives unemployment benefits at the point of exit, and zero otherwise. Exit year fixed effects are included. Age fixed effects by broad age group (20-29, 30-39, 40-49, 50-59, 60-69 and 70+), as well as income and wealth decile fixed effects are included. Observables are measured at the point of exit. Standard errors are clustered at the individual level. The regression uses data on exiters from 1994-2014.

3.4 Fact 4: re-entry often occurs soon after exit

Fact 3 established that re-entry does occur for a non-negligible proportion of participants. We now ask: conditional on occurring, how soon after exit do individuals re-enter? Figure 6 plots a histogram of the re-entry times observed in the data. Almost half of all re-entry occurs just 1 year after exit, indicating that re-entry tends to be quick. Combined with the evidence for short spells given in Section 3.1, this implies that there is a high degree of turnover between participation and non-participation states with many individuals dropping out of participation spells after only a few years and a non-negligible number re-entering soon after exit.

FIGURE 5: Impact of income, wealth and age on the probability of re-entry

(A) Income

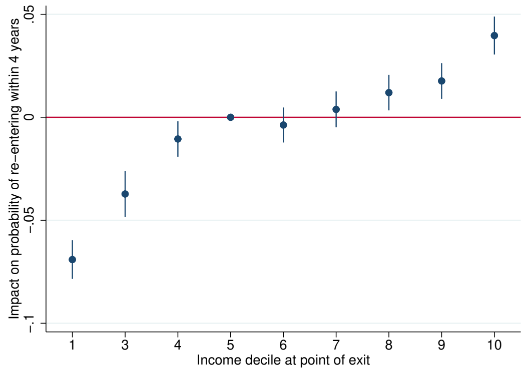

(B) Wealth

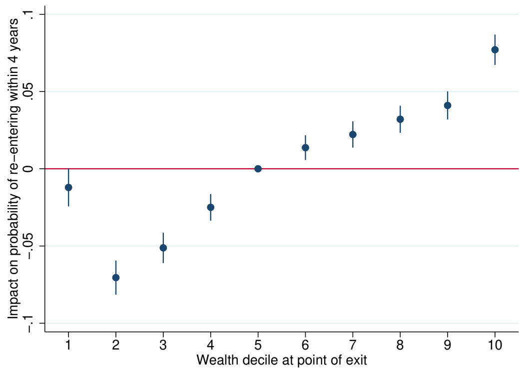

(C) Age

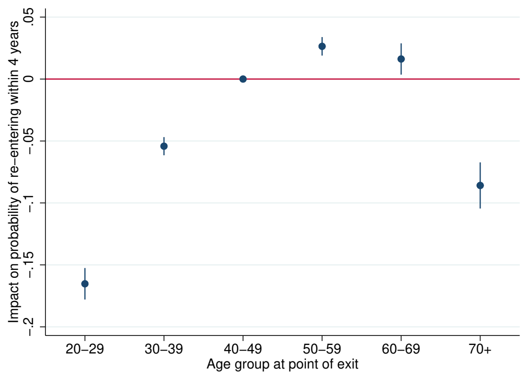

Note: this figure plots the coefficient estimates for the fixed effects on income and wealth deciles and age groups following estimation of Equation 2. This specification includes individual fixed effects. Variables are measured at the point of exit, and deciles are based on the full Norwegian population aged 20 and above in that year. Panel (A) shows the impact of income and panel (B) shows the impact of wealth. The effects are estimated relative to the median (5$^{\text{th}}$ decile). Panel (C) gives the effects of age. 95% confidence intervals are shown. The red line represents a null relative effect.

We apply similar robustness checks to those undertaken in Section 3.1 to verify this empirical fact: first, we check to see whether this quick re-entry could be driven by the receiving of gifts or inheritances. We apply the same three checks here, namely excluding individuals who receive a gift or inheritance above 10,000 NOK (≈ \$1670) in the year of or before re-entry, excluding re-entrants for whom a parent or grandparent died in the year of or before re-entry and removing re-entrants for whom a parent or grandparent held public equity in the year of or before re-entry. Figure D.10 gives the histogram for these subsamples and shows very

FIGURE 6: Distribution of re-entry times

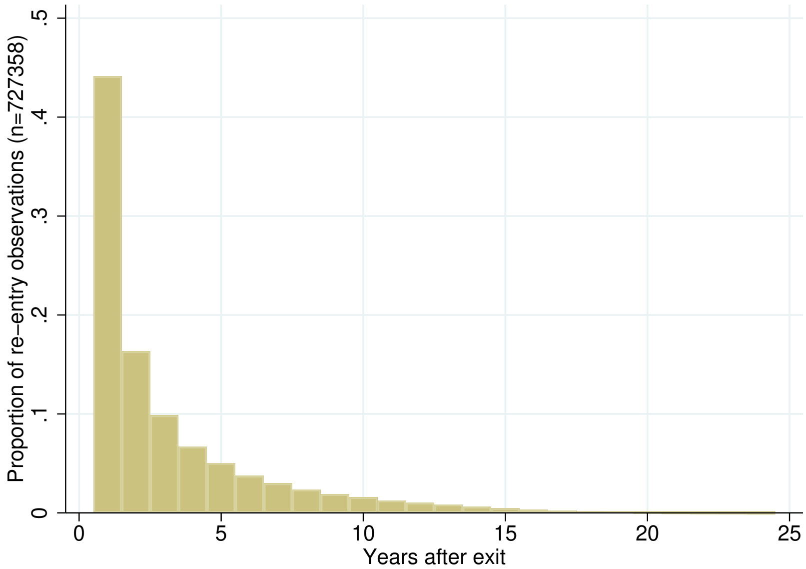

Note: this histogram plots the distribution of re-entry times in the Norwegian data. The x-axis gives the re-entry time (in years) and the y-axis shows the proportion of re-entry observations belonging to a particular length.

similar re-entry time distributions to the baseline figure. Second, we use the Shareholder Registry (available from 2004) to identify individuals who hold stocks in the company they work for when they re-enter. Figure D.11 plots the histogram excluding these re-entrants and shows very similar patterns.

3.5 Fact 5: downward-sloping hazard function for re-entry

Our final empirical fact studies how the likelihood of re-entry changes with the duration since exit. Our object of interest is $h_i(d)$, which gives the probability of re-entering the stock market d years after exiting conditional on not having re-entered until then. To do this, we take advantage of the fact that some individuals have multiple spells out of the stock market and apply the GMM estimator developed by Alvarez et al. (2021).

Figure 7 plots the estimated hazard function for re-entry. The hazard function is downward sloping and highly convex, indicating negative duration dependence in re-entry following exit: the longer it has been since you left the stock market, the lower is the probability of returning. There is a sharp decline in the hazard rate in the initial years following exit with the hazard rate at $d = 2$ being less than half that of $d = 1$. By $d = 12$, the hazard rate is very low, indicating that the likelihood of re-entering a decade after exit is virtually zero.

FIGURE 7: Baseline hazard function for re-entry

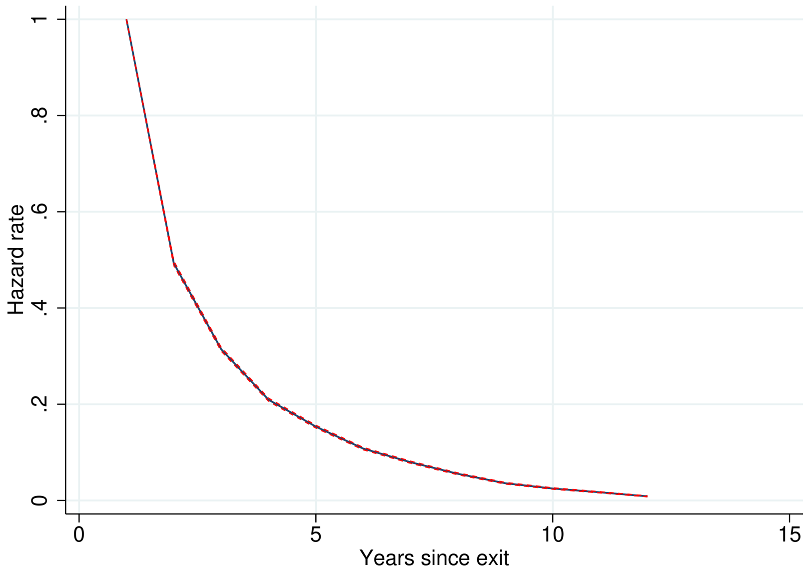

Note: this figure plots the estimated baseline hazard for re-entry following exit using the methodology of Alvarez et al. (2021) described in Section C.1. The dotted red lines denote 95% confidence intervals. The hazard rate at duration $d = 1$ is normalised to 1.

4 Are standard models of participation consistent with these dynamics?

The basic portfolio choice model à la Merton (1969) predicts that all individuals should participate in the stock market as long as the expected risk premium is positive. Given that the historical average risk premium in the stock market is well above zero, we should see participation rates of 100% at all points in time - and thus no dynamics - according to this simple rational model. However, the empirical finding that participation rates are far below 100% has generated a literature that provides a plethora of explanations. In this section, we examine broad categories of proposed explanations for this underparticipation puzzle and ask whether they can generate the dynamics observed in the Norwegian data. As discussed in

Gomes et al. (2021), explanations for underparticipation can be divided into four broad categories: non-standard preferences, participation costs, risks faced by households and social environments. We take each in turn and also consider alternative candidate explanations such as pension holdings and liquidity shocks in Section 4.5.

4.1 Non-standard preferences

Expected utility maximisers with standard preferences exhibiting second-order risk aversion (e.g. CRRA utility) should always be willing to invest some money in stocks as long as the expected risk premium is positive (Haliassos and Bertaut (1995)). This is because such individuals are effectively risk neutral for small risks and risk has no first-order effect. However, first-order risk aversion, whereby individuals have a kink in the utility function at some certainty point, can make risk aversion locally infinite and zero stockholdings an optimal outcome (Segal and Spivak (1990)). A range of preferences exist that exhibit first-order risk aversion including, but not limited to, prospect theory (Kahneman and Tversky (1979)), disappointment aversion (Ang et al. (2005)), news utility (Pagel (2018)) and ambiguity aversion (Cao et al. (2005)).

To generate dynamics in stock market participation, individuals need to display time-varying preferences. If preferences were time-invariant, then the set of non-participants and participants would be the same over time and there would be no movements in and out of the stock market. Note that time-varying preferences cannot simply be a change in the coefficient of relative risk aversion in CRRA utility as this would just change the optimal risky asset share, but not bring it down to zero. Instead, individuals need to switch between orders of risk aversion. In particular, to explain short spells, people need to switch from exhibiting second-order to first-order risk aversion soon after entry. To obtain quick re-entry, the opposite is required: some exiters need to switch back to exhibiting second-order risk aversion soon after exit. As such, not only is time-varying preferences needed to rationalise the facts, but we also need changes in preferences to occur at a fairly high frequency. The downward sloping hazard functions also mean the propensity to change preferences must fall with time spent in/out of the market.

In economic models, there is often the assumption that preferences are given and stable over time, and so changes in behaviour reflect changes in opportunity sets rather than preferences (Stigler and Becker (1977)). Empirical studies have typically found positive and significant correlations in individuals' risk preferences, though correlations are usually below 1 (Chuang and Schechter (2015); Dohmen et al. (2016)). As such, these correlations indicate that preferences are moderately stable; however, correlations are not perfect and so preferences may exhibit some movements over time.[23] Schildberg-Hörisch (2018) provides a framework for studying why risk aversion may change: first, individuals could become more risk-averse over the life cycle. This mechanism is consistent with empirical evidence in Dohmen et al. (2017), who find that willingness to take risk falls linearly until age 65, after which the slope flattens.[24] However, the dynamics in stock market participation we observe are at a high frequency, whereas age-induced movements would occur slowly.

Second, economic crises and downturns could lead to shifts in risk aversion. For example, there is evidence that risk aversion has increased since the financial crisis (Dohmen et al. (2016); Guiso et al. (2018)). Furthermore, Malmendier and Nagel (2011) show that willingness to take financial risks depends on the aggregate stock market returns over the course of an individual's life. One could therefore argue that short spells could be linked to poor aggregate stock market performance that causes individuals to switch from second- to first-order risk averse preferences. Figure D.12 plots the proportion of entrants in each given year who exit within 1, 2 or 3 years after entry. We do not observe isolated jumps in the prevalence of short spells around periods with aggregate stock market downturns, namely the bursting of the dot-com bubble in the early 2000s or the financial crisis. Instead, we see that the degree of short spelling steadily grew from the mid-1990s until around 2003. There is no distinct increase in the early 2000s. Furthermore, the proportion remains fairly steady around the financial crisis. As such, our finding of short spells is not driven by periods of aggregate stock market downturns.[25]

[23]It is noteworthy that part of the imperfect correlations observed in panel data studies could reflect measurement error (Schildberg-Hörisch (2018)).

[24]Schurer (2015) also finds a decline in risk tolerance up to age 45. Beyond this age, changes in risk tolerance depend on socioeconomic status. Other papers that have found a link between age and risk aversion include Levin et al. (2007) and Paulsen et al. (2012).

[25]One could argue that individual-level experiences rather than macroeconomic conditions are what matter. However, Sahm (2012) find that individual-level events such as changes in income or wealth, job displacement or being diagnosed with a serious illness have little effect on risk tolerance. Instead, there is a role for macroeconomic conditions.

Third, temporary swings in risk aversion could be induced by stress, fear or other related emotions. Papers have shown that negative emotions increase risk aversion (e.g. Kandasamy et al. (2014); Cohn et al. (2015); Guiso et al. (2018)).[26] While we cannot completely rule this out as an explanation for our findings, we note that Schildberg-Hörisch (2018) states that these factors should generate typically small changes in risk preferences. In our setting, we require not that people sometimes feel more or less risk averse, but that they completely switch the order of risk aversion reflected in their preferences. As such, we require these emotions to trigger a significant change in preferences.

4.2 Participation costs

A leading explanation for limited stock market participation is participation costs (Haliassos and Bertaut (1995); Vissing-Jørgensen (2002); Gomes and Michaelides (2005)). Such costs can reflect direct monetary expenditures associated with investing (e.g. fees for setting up a brokerage account), as well as pecuniary and informational costs. Vissing-Jørgensen (2002) gives evidence in support of fixed participation costs. She considers two types of fixed costs: the first is a fixed transactions cost, which can reflect time spent implementing trades and, in the case of first-time buyers, time spent acquiring knowledge of fundamental investment principles. Such costs are effectively entry costs for non-participants and exit costs for participants. They thus provide a cost to changing participation status, which can explain why some individuals remain non-participants. Vissing-Jørgensen (2002) finds support for such state dependence using PSID data. The second type of fixed cost is a per-period participation cost that, for example, can capture time spent monitoring your accounts over the year. Vissing-Jørgensen (2002) estimates that a per-period cost of just $50 per year (in year 2000 prices) can explain why half of non-participants choose not to participate.[27]

While fairly small fixed costs are sufficient to explain why many individuals do not participate, our question is whether such costs can generate the dynamics we observe. Fixed transaction costs should slow down dynamics of participation because it is costly to quickly exit only to re-enter soon after, which conflicts with our empirical findings. In principle,

[26] This finding fits with the Affect Infusion Model of Forgas (1995), which predicts that people in a bad mood should be more risk averse as they become more aware of downside risks.

[27] Vissing-Jørgensen (2002) also considers variable transaction costs, whereby the cost of trading is directly proportional to the value of stocks bought/sold; however, she does not find support for such costs in the data.

participation costs, whether per-period costs of participating or fixed transaction costs, can generate entry and exit through fluctuations in financial wealth or income that make such costs binding. Entry can occur if non-participants experience a rise in their investable wealth, though to obtain short spells, we would need this wealth to fall again soon after entry. Quick re-entry would also require wealth to rise again soon after exit. As such, to obtain the dynamics we observe in the data, we would require a very volatile process for investable wealth. A further challenge for the participation cost story, even for explaining the static underparticipation puzzle, is explaining why wealthy individuals exit as fixed costs should not influence their participation decision. Indeed, empirical studies have found participation rates to be below 100% even for the wealthiest households (Guiso and Sodini (2013)). While we do see that short spells are relatively more prevalent for the low income and wealth groups (Figures 2a and 2b), Figures D.9a and D.9b show that short spells are still very prevalent for those with high income and wealth respectively.

4.3 Risks faced by households

A strand of the literature has studied how "background risks", particularly labour income risk, can affect portfolio allocations. Theoretically, the impact of labour income risk depends on the nature of the risk (Vissing-Jørgensen (2002)): first, if labour income is riskless, this should lead to a higher investment in risky financial assets because in effect, such labour income is equivalent to holding a riskless bond. Second, if labour income is risky but uncorrelated with stock returns, then you should tilt your portfolio away from stocks as there is already risk coming from your human wealth.[28] Third, if labour income is risky and correlated with stock returns, then there is a hedging component that runs in the opposite sign of the correlation. For example, if business cycle risk produces a positive correlation between labour income and stock returns, then the optimal portfolio choice requires you to reduce stockholdings (Haliassos and Bertaut (1995)). It is important to note that zero stockholding cannot be an optimal solution in the first two cases. Risky labour income that is uncorre-

[28] Fagereng et al. (2017b) studies the impact of uninsurable wage risk on portfolio shares using Norwegian data. They find a significant marginal effect of such risk on portfolio shares, although the economic impact is limited because the size of this wage risk is small. Vissing-Jørgensen (2002) finds a negative impact of the volatility of non-financial income on both the probability of stock market participation and the proportion of wealth invested in stocks conditional on participating.

lated with stock returns serves to reduce the optimal portfolio share, but would not push it down to zero. However, Haliassos and Bertaut (1995) show that, particularly if coupled with a no short selling constraint, zero stockholding can be an optimal choice for sufficiently low wealth if there is a positive correlation between labour income and stock returns. Therefore, to generate entry and exit, it is not sufficient for the level of wage risk to change over time. Instead, we require the correlation between labour income and stock returns to be time-varying, which could be harder to justify. Furthermore, if households exhibit standard CRRA preferences, then the hedging motive should not lead to zero stockholding for high wealth groups. This is because with CRRA utility, those with high wealth care less about insuring against bad states and so would continue to participate in the stock market. However, as shown in Figure D.9b, short spells do still occur for high wealth individuals.

4.4 Cultural and social environment

Cultural factors can influence an individual's beliefs and preferences, which in turn can affect economic outcomes (Guiso et al. (2006)).[29] Various papers have provided empirical evidence of a causal link running from cultural environments to savings behaviour, often by studying immigrants of different cultures who move to a common country and thus face the same institutional and policy environment. Haliassos et al. (2017) study migrants to Sweden and find significant differences in financial behaviour and the propensity to hold stocks based on the degree of cultural similarity to Sweden.[30] While underparticipation in the stock market could be linked to cultural factors, these factors need to be time-varying in order to obtain dynamics in participation. However, Guiso et al. (2006) define culture as *"customary beliefs and values that ethnic, religious and social groups transmit fairly unchanged from generation to generation"*. As such, the view is that cultural factors that come from, for example, religion and ethnic background are very slow-moving and thus would not be able to reproduce the high frequency entry and exit that we observe.

Instead, social interactions could generate more frequent changes in beliefs and preferences. Shiller et al. (1984) argues that investing is a social activity and so investment de-

[29] For example, ethnic origin has been shown to affect trust (Guiso et al. (2003)).

[30] Other papers that find significant effects of culture of financial behaviour include Osili and Paulson (2008), Guin (2017) and Fuchs-Schündeln et al. (2020). However, some papers do not find such effects (Carroll et al. (1994, 1999)).

cisions can be affected by the actions of those you interact with. A growing literature has given empirical evidence on the influence of peer effects on financial behaviour.[31] In principle, communication between peers could lead to entry and exit. If my neighbour decides to leave the stock market - perhaps due to experiencing poor returns - this could induce me to also leave. However, we argue that peer effects will struggle to explain all of the dynamics we observe for a variety of reasons: first, Kaustia and Knüpfer (2012) show that good stock returns experienced by local peers can positively affect an individual's decision to enter the stock market. However, the authors do not find evidence of a discouragement effect following poor realisations, from which they infer that peers primarily share good outcomes. Therefore, peer effects could struggle to explain exit, particularly the quick exit we found in Fact 1. Second, it is difficult to rationalise the downward-sloping hazard functions obtained in Facts 2 and 5 as these imply that the effect of peers diminishes with time. Third, our focus is on the extensive margin of participation and so we require social interactions to generate complete exit rather than just exit from a particular stock. One could imagine individuals discussing particular stocks and perhaps a bad return experienced by a peer may deter you from also investing in that security; however, it may not necessarily put someone off investing in other stocks.

4.5 Other candidate explanations

4.5.1 Liquidity shocks

In principle, individuals might have to leave the stock market due to liquidity needs. For example, people may lose their job or face unexpected health expenses. Upon the "completion" of such liquidity needs, individuals may subsequently re-enter the market. In general, one might expect a constant Poisson arrival of such shocks. However, a constant arrival rate would imply a flat hazard function of exit from participation, which contradicts the downward-sloping hazard function estimated in Figure 3. For liquidity shocks to therefore be consistent with our hazard function, we would need these shocks to be more likely to oc-

[31] Hong et al. (2004) show that households who report interacting with their neighbours and attending church are more likely to participate in the stock market even after controlling for individual characteristics and personality traits. Brown et al. (2008) find a causal link between individual stockholding and the average participation of the individual's community, which they argues occurs through word-of-mouth communication.

cur early in the spell. A priori, there is no clear reason why this should be the case. Indeed, if anything one might expect the reverse as people would likely not enter the stock market in the year before any expected liquidity needs such as a house purchase given the risk of a stock market downturn. Therefore, the nature of the hazard function suggests that liquidity shocks are not the driver. In order to further verify this, we link to other administrative datasets to investigate whether the propensity of observable liquidity shocks varies by spell length. In particular, we look at house purchases, divorce and unemployment as our liquidity shocks.[32] Figure D.14 plots the proportion of exiters of different spell lengths experiencing at least one of these three shocks in their exit year. For comparison, we also show the proportion of non-exiters (both non-participants and continuing participants combined) experiencing a liquidity shock. We can see that some exit is correlated with such shocks: about 7% of non-exiters experience a liquidity shock compared to around 11-12% for exiters. Indeed, this is in line with the literature which shows that exit can be linked to house purchases (Brandsaas (2021)), marital status (Christiansen et al. (2015)) and unemployment (Basten et al. (2016)). However, the prevalence of liquidity shocks is very similar across spell lengths, suggesting that short spellers do not have a higher likelihood of facing a liquidity shock compared to longer spellers. Furthermore, if around 12% of exiters leave because of one of these observed shocks, it means that 88% of exiters are leaving for other reasons. All together, it appears that liquidity shocks are unlikely to explain the prevalence of short and multiple spells in the stock market.

4.5.2 Sophisticated market timing

Could the short-lived entry and exit observed in the data be driven by sophisticated market timers? Perhaps these individuals pursue short-term investment strategies and re-enter whenever a promising investment opportunity arises. If this were the case, we would expect short spelling to be correlated with proxies for financial sophistication. However, Table 2 and Figure 2 show that short spelling is negatively correlated with characteristics typically

[32]Two other liquidity needs could be health shocks and education costs (perhaps for children). However, higher education is free in Norway. While healthcare is not free, there is an annual deductible above which healthcare is free. This deductible is fairly small at NOK 2,460 in 2021 ($410 in 2011 USD). Across OECD countries, Norway had the highest share of healthcare financed through government schemes and the largest per capita spending on healthcare relating to long-term care (Cooper (2019)). As such, Norwegians in general seem not to be susceptible to high financial costs linked to healthcare needs.

associated with higher financial literacy (college education, income and wealth). Furthermore, we would expect sophisticated market timers to do better than other participants. Figure 8 looks at the performance of exiters of different spell lengths. Here, we measure performance by computing the proportion of exiters of different spell lengths reporting only taxable gains from the sale of stocks and equity funds (Figure 8a) or only losses (Figure 8b) in their exit year.[33] Short spellers of 1-2 years are less likely to report only gains and more likely to report only losses. The unconditional probability of reporting only gains is 30% for short spellers compared to around 40% for those participating for longer. Similarly, the unconditional probability of reporting only losses for short spellers (≈28%) is twice that of longer spellers (≈13-15%). Taken together, these figures suggest that short spellers have weaker average performance compared to longer spellers.[34]

4.5.3 Pensions

One may worry that the existence of pension wealth could affect individuals' desire to actively invest in the stock market out of their non-pension wealth. In principle, a rational agent should consider their overall portfolio, comprising of both pension and non-pension wealth, together when deciding upon their optimal portfolio allocation. If, for example, your pension wealth is already invested in the stock market, you may invest less (or nothing at all) out of your remaining wealth. As such, observing non-participation in the tax records, which do not contain data on occupational or public pension wealth, may not necessarily mean that an individual has no exposure to the stock market. Non-participation out of non-pension wealth could simply be a rational choice given existing exposure through pensions.

To be able to explain dynamics, it would need to be the case that: 1) the desired risky asset share out of *total* wealth changes and individuals adjust their non-pension holdings to achieve this new goal, and/or 2) exposure to the stock market coming from pension wealth

[33] For this analysis, we restrict attention to exiters who entered from 2006 onwards because of changes in the Norwegian tax system that can make it difficult to interpret the tax record variables prior to this point. From 2006 onwards, individuals were only taxed on capital gains above a risk-free return. However, before 2006 the taxable amount depended on the share's proportion of retained taxed capital, and so may not necessarily be linked to achieving a high/low return relative to a risk-free aseet. Taxed capital refers to undistributed income that has been previously subject to tax at the company level. Focusing only on exiters who entered from 2006 onwards aids with the interpretation of the tax variables because these individuals would be subject to the "new" tax system based on risk-free deductions.

[34] Figure D.13 plots the corresponding figures based on reporting any gains or any losses rather than only gains or losses. We obtain broadly similar findings.

FIGURE 8: Performance of exiters by spell length

(A) Report only gains

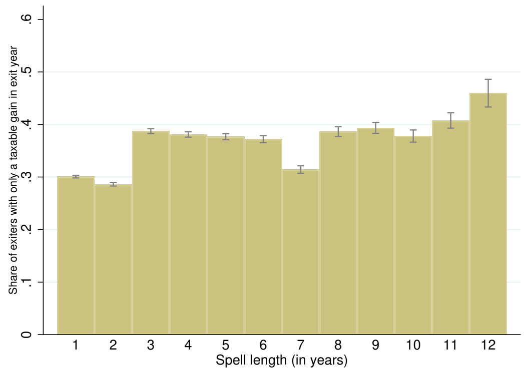

(B) Report only losses

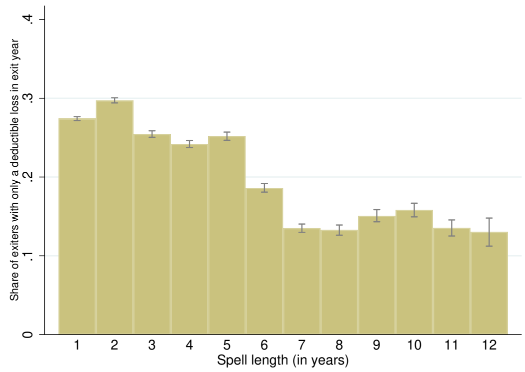

Note: this figure shows the performance of exiters by spell length based on records of taxable gains and tax-deductible losses in the income tax data. In panel (A), we plot the proportion of exiters of a given spell length reporting only gains from the sale of stocks and funds (computed as the sum of items TR 3.1.8, TR 3.1.9 and TR 3.1.10 in the tax records) in their exit year. In panel (B), we plot the proportion of exiters reporting only losses (computed as the sum of items TR 3.3.8, TR 3.3.9 and TR 3.3.10). We use exiters who enter from 2006 onwards in these plots.

is changing at a high frequency and individuals identify these changes and adjust their portfolio accordingly. Explaining frequent exit and (re-)entry through this rebalancing channel is arguably difficult as it requires individuals to regularly follow movements in their pension holdings and to actively rebalance accordingly. Various papers have shown that portfolio adjustments in retirement accounts are sluggish using data on 401(k) retirement accounts in the US (e.g. Agnew et al. (2003); Ameriks and Zeldes (2004)). Indeed, there is also evidence outside of retirement accounts that investors can be slow to rebalance (e.g. Brunnermeier and Nagel (2008); Calvet et al. (2009a); Karlsson et al. (2009)).[35] As such, the evidence suggests that portfolio rebalancing involving both non-pension and pension accounts combined is unlikely to occur at the high frequency required to explain quick exit and re-entry.

Nevertheless, we undertake a reading of the Norwegian pension system to better understand whether the nature of the system could interact with our empirical findings. There are three main components: first, the National Insurance Scheme (*"folketrygden"*) is the basic public pension scheme in Norway and ensures everyone receives a minimum pension

[35] Different explanations have been proposed for limited active rebalancing including bounded rationality (Sims (2003)), observation costs (Abel et al. (2007, 2013) and news utility (Pagel (2018)).

income. Furthermore, workers are guaranteed a supplement that is proportional to their income.[36] A key feature of the public pensions system is that it is a defined-benefit system and so citizens face no stock market exposure through this. As such, the decisions to exit and enter the stock market cannot be attributed to portfolio rebalancing between private accounts and public pension wealth.

Second, there are occupational pensions. Public occupational pensions are also defined-benefit in nature, meaning no exposure to stock market risk through such schemes.[37] Private sector occupational pensions operate differently. Until 2001, only defined-benefit pensions existed. While defined-contribution pensions, for which the pension benefit depends on how well the contributions are invested, were allowed from 2001, they did not gain momentum until 2006 when occupational pensions were made mandatory by law. Indeed, before 2006 occupational pensions were mainly provided by larger employers (OECD (2009)).[38] A concern could therefore be that private sector defined-contribution occupational pensions have some exposure to the stock market and this could influence choices made in non-retirement investment accounts. However, this exposure really begins only from 2006 and if it were interacting with stock market participation through private accounts, then we should see greater dynamics following the rise of such pensions. However, Figure D.12 shows that the prevalence of short spells was highest in the early 2000s and there is no drastic jump after 2006.

Third, individuals may have personal private pensions that they invest into. As payments into an Individual Pension Scheme (IPS) in Norway are tax deductible up to a certain limit, one can infer from the tax records whether an individual holds such a scheme.[39] Figure D.15

[36] Under the current system, in each year of employment 18.1% of your wages up to a certain ceiling is transferred to your pension account. This pension income is then indexed to nominal wage growth. Upon retirement, the accumulated amount is not given as a lump sum. Instead, an annual sum is given based on the expected number of years you will be a pensioner, which itself depends on when you first start withdrawing your pension and life expectancy. While there are some differences based on your year of birth, the overall premise of pensionable income being linked to your employment earnings still holds. For further details, see Fagereng et al. (2019) and Fredriksen and Halvorsen (2019).

[37] Until 2020, the public occupational pension scheme was such that workers were entitled to the maximum pension after 30 years of service and can get a pension equal to 66% of their pension base (equal to their final salary converted into a full-time equivalent) before adjustments for life expectancy. However, from 2020 occupational pension earnings became similar to that in the National Insurance Scheme, in particular having a share of your earnings each year be accumulated in a pension pot. However, this remained a defined-benefit system. For further details on public occupational pensions and the reforms, see Fredriksen and Stølen (2018).

[38] As of 2018, 90% of private sector employees were under a defined-contribution pension (Fredriksen and Halvorsen (2019)).

[39] There are two relevant variables in the tax data. Item 3.3.5 records the deductible amount from payments

provides a time series of participation in private pension accounts separately for the whole population and the subset of the population aged 60 or under (who are unlikely to have drawn from such pensions yet). In either case, the participation rates are in single digits, indicating that the vast majority of the population do not hold private pension accounts. To further ease concerns that interactions with private pensions are unlikely to drive our results, we plot the proportion of exiters of different spell lengths who hold private pensions as of their exit year. If these schemes were driving our short spell result, we might expect to see a greater prevalence of private pensions amongst short spellers; however, Figure D.16 shows the opposite. We also reproduce our spell length histogram, but excluding any individual who at any point in the sample holds a private pension account. Figure D.17 shows that our results are robust to this. We therefore believe that pension holdings cannot explain the dynamics we observe.

4.5.4 Tax optimisation

Could the quick exit and re-entry from the stock market be due to tax optimisation? Perhaps individuals choose to exit to reduce their tax liability in a given year. There are two possible tax margins that could be a cause of concern. The first is the wealth tax, whereby individuals are taxed on net wealth above a given threshold.[40] However, most individuals do not reach the threshold particularly because the tax value on housing is 25% of its market value. As such, it is very unlikely that a desire to avoid the wealth tax can explain the entry and exit decisions of Norwegian individuals given that most of them will not be subject to the tax. Indeed, stocks and mutual fund holdings are given a valuation discount of 45% (in 2021), whereas cash or deposit account holdings are not given a discount, and so it is actually better for wealth tax purposes to retain wealth in stocks and funds rather than liquidating. The second relevant tax is capital gains tax. In Norway, losses made from the sale of stocks and equity funds are tax deductible, while gains above a risk-free return are taxed at an effective rate of 31.68% (in 2021). Therefore, one might be worried that the quick exit we observe is

into an IPS, while item 4.5.1 indicates capital in an Individual Pension Account (IPA). Note that IPAs were replaced by the IPS in 2006, from which point new money could not be placed into your existing IPA and new IPAs could not be opened. We consider an individual to be a private pension contributor if they report a positive value for either of these two variables, either in the current year or in any past year.

[40] In 2021, net wealth above 1.5m NOK (≈$250,000 in 2011 USD) was taxed at 0.85% (0.7% to the municipality and 0.15% to the state). The threshold is doubled for couples.

because individuals are liquidating their loss-making shares to reduce their tax liabilities.[41] However, capital gains taxation in Norway is tied to the realisation for each individual security, not the performance of the overall portfolio. To be able to explain the complete exit that we observe, we would require every security in one's portfolio to be making a loss. Therefore, we argue that tax-motivated selling is unlikely to drive our results.

5 Model

In this section, we first evaluate the performance of the workhorse life-cycle model of Cocco et al. (2005) augmented with fixed participation costs in generating short-term dynamics. Upon showing that this model cannot produce these dynamics when calibrated to the Norwegian economy, we then augment an otherwise standard portfolio choice model à la Merton (1969) with three features of human behaviour established in existing work, and evaluate whether such a model can explain the dynamics observed in the data.

5.1 A life-cycle model with participation costs

The workhorse life-cycle model of Cocco et al. (2005) is a discrete-time model of consumption and portfolio choice. Investors face a finite horizon T that is divided into two subperiods: working age ($t \leq t_r$) and retirement ($t > t_r$). During working life, labour income is subject to undiversifiable shocks. In particular, labour income has a deterministic age component but is hit with both permanent and transitory shocks. In retirement, labour income is constant. Investors face borrowing and short-selling constraints, and can invest in two financial assets: a risky asset (stocks) and a riskless asset (bonds). The only modification made to the Cocco et al. (2005) model is the inclusion of fixed participation costs. In the absence of participation costs, individuals would invest at least a small amount in the stock market in every period other than at age T given that the expected risk premium on stocks is positive.[42] This follows from the standard Merton rule (Merton (1969)). The inclusion of fixed participation costs gives a reason to not participate - if investable wealth is low rela-

[41] This relates to findings in Odean (1998), who shows that the prevalence to sell losing stocks is highest in December, which can be linked to the end of the tax year and attempts to reduce tax liability.

[42] At age T, you know you will die in the next period and so optimally choose not to save.

tive to the participation costs, then it may be optimal for individuals not to participate.[43] Following Gomes and Smirnova (2021), we include both on-entry costs and per-period participation costs. The full model setup and calibrated parameter values are given in Section B.1. In theory, such a model with participation costs could generate some short spells and re-entry. This is because some people may receive a bad income or return shock that reduces their cash on hand the following period. As a result, their investable wealth is lower and they decide it is no longer worthwhile to pay the participation costs to continue participating. As they build up more wealth (or receive a positive income shock), they may decide to re-enter. Ultimately, this becomes a quantitative question as it will depend on, amongst other factors, the magnitude of shocks and the size of the participation costs. Therefore, the objective of this section is to understand whether we can obtain the dynamics observed in the Norwegian data in a standard life-cycle model calibrated to the Norwegian economy.

We simulate the model for $N = 10,000$ individuals.[44] Before examining the performance of the model in generating the quick exit and re-entry dynamics, we first analyse the life-cycle patterns for participation and risky asset share (conditional on participating). Figure 9 plots the simulated participation rate over age. We see that participation is not 100% at all ages unlike in the workhorse Cocco et al. (2005) model. This reflects the inclusion of participation costs which can make it optimal to not participate. The participation rate rises sharply during early life. By age 30, everyone is in the stock market and the participation rate remains at 100% until early retirement when we start seeing exit. Exit occurs in retirement because individuals have a lower income and start decumulating the wealth they have built up during working life in preparation for retirement. As such, your investable wealth is decreasing as you get older during retirement, which means that eventually it is not worthwhile paying the participation cost, resulting in exit. The participation rate then falls sharply as you approach death.

Figure 10 plots the conditional risky asset share. We find a similar pattern to that of the workhorse model. First consider retirement age - the portfolio share in stocks increases with age (other than for the very final years before death). The intuition for this follows

[43]This model can be seen as a simplified version of Gomes and Smirnova (2021), who also augment the workhorse Cocco et al. (2005) model with fixed participation costs. However, they have further modifications that we do not consider in our setup such as unemployment shocks, consumption floors and heterogeneous savings motives across individuals.

[44]In the simulation, individuals are born with zero wealth.

FIGURE 9: Simulated participation rates

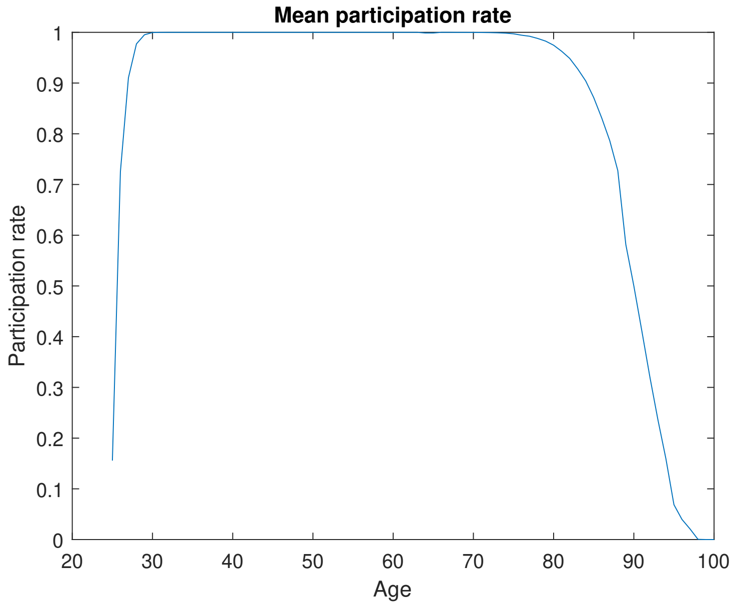

Note: this figure plots the simulated participation rate over age based on a simulation of 10,000 individuals.

from Jagannathan and Kocherlakota (1996), Cocco et al. (2005) and Gomes (2020). During retirement, agents receive a non-random stream of labour income. This can be thought of as an endowment of the riskless asset. As you get older, you are decumulating your wealth so the ratio of this income endowment to wealth is going up. Without any active changes to your portfolio allocation, the risky asset share out of *total* wealth (both labour income and financial) is going down. Therefore, individuals want to tilt their financial portfolio more into stocks, which pushes up the conditional risky asset share. During working age, the reverse occurs - as you get older, you reduce your risky asset share. This is because agents are accumulating a lot of wealth during this time while labour income is not increasing as fast (see Figure D.18), so the ratio of labour income to wealth is going down with age. As such, the implicit risky asset share is high and so agents should tilt their financial portfolio more towards the riskless asset.[45]

Moving onto the dynamics of participation, Figure 11 shows the distribution of spell

[45]Note that labour income is not entirely riskless in this period due to the permanent and transitory shocks. However, one can think of the correlation between labour income shocks and stock returns as being low (see Davis and Willen (2013); Gomes (2020)). In this calibration, the correlation is set to zero. As such, labour income is still a close substitute to the riskless bond.

FIGURE 10: Simulated mean conditional portfolio share in stocks

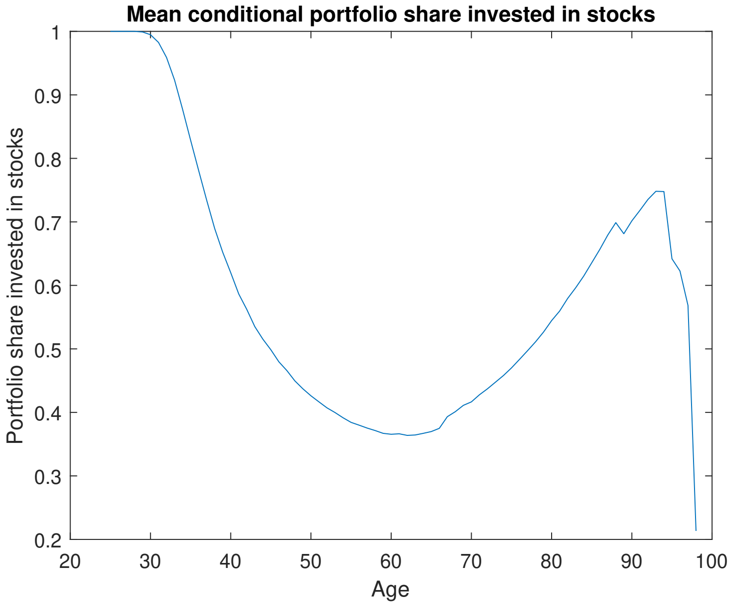

Note: this figure plots the simulated mean portfolio share invested in stocks conditional on participating in the stock market. This is based on a simulation of 10,000 individuals.

lengths in the simulated life-cycle model. We see that the model only generates long spells with the earliest exit time being 29 years after entry. Most individuals exit as they approach death and are decumulating their wealth holdings. The model is therefore not capable of explaining short-term participation fluctuations given these parameter values. Figure D.19 plots the hazard function for exit from participation. In line with the histogram of spell lengths, the hazard rate is virtually zero until a spell length of 50 years after which the hazard rate increases towards 1. The increase towards 1 is simply because individuals never participate in the final year of life and so would certainly exit at that age if they had not exited already.

Moving onto the re-entry facts, Figure 12 looks at how many spells people have over their life in the simulated model compared against the Norwegian data. In line with Figure 9 which shows that the participation rate is 100% for much of working life, there are no individuals who never participate in the model. This contradicts with the data which shows a "never-participant" share of 43%. Instead effectively everyone has a single spell. 0.25% of individuals in the simulated model had more than 1 spell, far less than the 17% in the Norwegian data. It is difficult to analyse re-entry times in this model because we obtain such

FIGURE 11: Simulated spell lengths

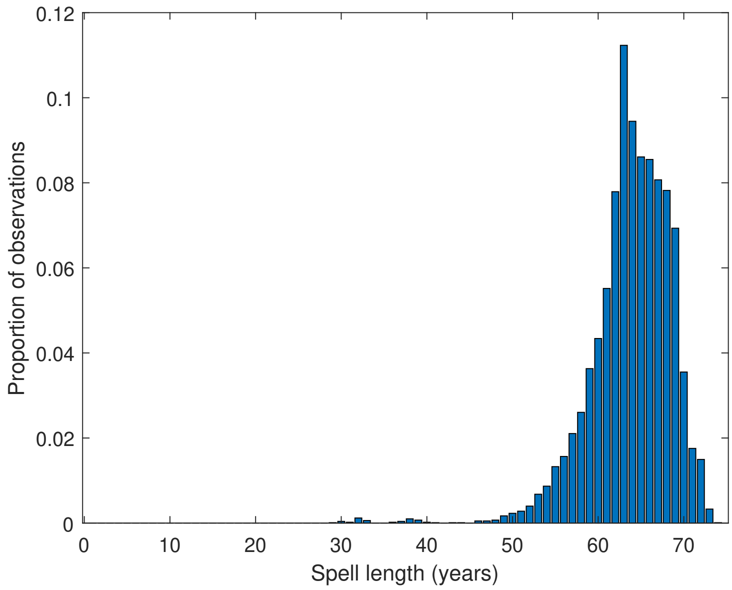

Note: this figure plots a histogram of the simulated spell lengths. This is based on a simulation of 10,000 individuals.

little re-entry (just 25 individuals out of 10,000 simulated individuals re-enter). Nevertheless, we plot the distribution of the re-entry times we observe in Figure D.20. Virtually every re-entry time is of just 1 year in the simulated model. Given that there is such little re-entry and conditional on occurring it happens almost certainly just 1 year after exit, there is no meaningful re-entry hazard function.

5.2 Three characteristics of human behaviour

Before giving the details of our proposed model, we first discuss three characteristics of human behaviour. The first characteristic is *heterogeneous ability* in the stock market: some individuals are able to generate higher returns on average than others, perhaps due to education or talent differences. Gabaix et al. (2016) propose this "type dependence" as a potential mechanism that can generate a positive correlation between wealth and returns. High ability individuals can earn persistently higher returns, allowing them to accumulate more wealth and reach the top of the wealth distribution. Empirical support for this mechanism is given by Fagereng et al. (2020), who use Norwegian data and find that including individual fixed effects to capture persistent heterogeneity can increase the explained variability in

FIGURE 12: Simulated number of spells

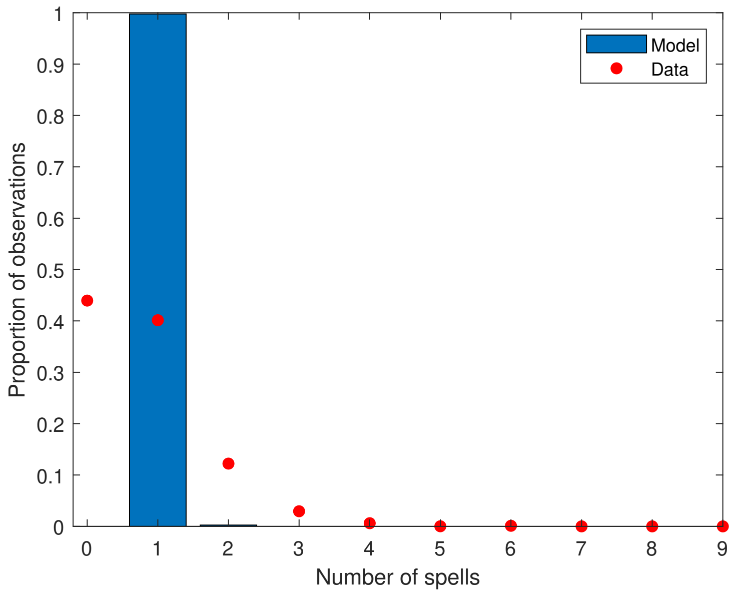

Note: this figure plots the distribution of the number of spells the simulated individuals experience over their lifetime. This is based on a simulation of 10,000 individuals.

returns from one-third to one-half. Bach et al. (2020) use Swedish data and show that type dependence does contribute to the heterogeneity of wealth returns. Other papers have also provided evidence of a positive correlation between financial literacy and the return to investments (see, amongst others, Gaudecker (2015); Bianchi (2018); Deuflhard et al. (2018)). In our model, we incorporate heterogeneous ability by having two types of individuals who earn different returns on average.

The second feature is that individuals have *incomplete information* about their ability types and thus need to *learn* about their ability. When individuals first enter into the stock market, they are unaware of how they will do. Indeed, the stock market is a complex environment and so no individual can perfectly know how they will perform. However, experienced returns give individuals a signal of their ability. Seru et al. (2010) find strong support for learning about ex-ante ability using Finnish transaction-level data. In particular, they find that most learning-by-trading occurs through individuals learning about their own ability and low ability individuals exiting the market. They also show that an investor whose performance is 1 standard deviation below the mean is about 15% less likely to continue trading, which suggests that retail investors do respond to their experienced returns. Linnainmaa

(2011) estimates a structural model and shows that investors trade to learn: if they have a successful trade, they infer skill and trade more, but following losses, they trade less. Given sufficient losses, investors will exit the market. Mahani and Bernhardt (2007) look at the effect of investor learning on market prices by embedding learning into a general equilibrium model. They show that learning reduces bid-ask spreads and the price impact of liquidity shocks. The idea of learning from experiences relates to the literature on experience effects (e.g. Malmendier and Tate (2005); Greenwood and Nagel (2009); Malmendier and Nagel (2011, 2015)). We capture this feature in the model by having individuals update their beliefs of being low ability over time based on their realised returns using standard Bayesian updating.

The third feature is that individuals have *noisy memory*, i.e. they cannot recall events perfectly. The notion of imperfect memory can be traced back to Ebbinghaus (1885), who shows empirically that memories decay and become noisier over time. A view in the psychology literature is that memories are costly to store and so it is difficult to remember past events precisely. Indeed, different memory systems are used for different types of information (see Poldrack and Foerde (2008)). Brocas and Carrillo (2016) propose a theory of optimal memory and show that extreme/exceptional experiences are stored using declarative memory, which is more accurate but more costly to store in. Azeredo da Silveira et al. (2020) study the optimal structure of memories and find that for a class of linear-quadratic-Gaussian forecasting problems, the optimal memory structure is one-dimensional, whereby individuals recall a single summary statistic (with noise) of their past experience. Therefore, in our setting, individuals do not remember their past stock market returns precisely. Instead, they recall a noisy version of their average annual return.

5.3 Model setup

Given that life-cycle patterns and participation costs only seem to help with explaining longer-term exit patterns based on the analysis in Section 5.1, we abstract away from life-cycle considerations and participation costs to simplify the model. There are T periods. In each period t, N_t new entrants enter the stock market for the first time. Each agent i can invest in a safe asset (bond) with a risk-free return r_s and a risky financial asset (stocks) with an idiosyncratic stochastic return r_{it}. Short-selling is not allowed in the model. Note that the return on the risky asset varies with i because each individual draws their own return, reflecting the fact that different individuals choose different stocks and funds to invest in and thus would have heterogeneous returns.

In order to capture the first feature of human behaviour, namely heterogeneous abilities, we assume that individuals are of one of two ability types: high (h) or low (l) ability. A share a_l of the N_t entrants in a given period are of low ability. Low ability individuals draw returns from a normal distribution with a lower mean than high ability individuals. In particular, an individual of ability type $y \in \{l, h\}$ draws return from a normal distribution with $r_{it}|y \sim N(\mu_y, \sigma^2)$. We assume *iid* draws conditional on type and $\mu_l < r_s < \mu_h$. The latter assumption means that if agents perfectly observed their type, the low ability individuals would never participate and the high ability individuals would always participate. This assumption therefore means it will not be the case that the expected risk premium is positive for all agents at all points in time, which means we can generate non-participation.

To capture the third feature of human behaviour (noisy memory), we assume two aspects of memory recollection: the first is that memory storage is costly, and so it is optimal for individuals to recall a summary statistic of their past experiences rather than each experienced return separately (Azeredo da Silveira et al. (2020)). Therefore, we assume that individuals recall the arithmetic average of their experienced return. The second is that experiences are recalled with noise. This can be interpreted in two ways: it is costly to store memories with perfect precision or alternatively it is difficult to precisely calculate your return and requires some financial expertise. Taken together, for an individual i at time t who participated in risky financial assets for $s \leq t$ periods, their recollection of their experienced returns is given by:

$$m_{it} = \frac{1}{s} \sum_{q=1}^{t} r_{iq} \cdot \mathbb{1}(\text{part}_{iq} = 1) + \epsilon_{it}$$

where $\epsilon_{it} \sim (0, \sigma_\epsilon^2)$, $\epsilon_{it} \perp r_{iq} \forall (t, q)$ and part_{iq} is a binary variable equal to 1 if individual i participates in risky financial assets in period q. The first term is thus the actual average return experienced by individual i as of period t. The second term reflects noise in the recollection of this average return. Given the assumptions of normality of ϵ_{it} and its orthogonality to returns r_{iq}, m_{it} also follows a conditional normal distribution with:

$$m_{it}|y \sim N(\mu_y, \frac{\sigma^2}{s} + \sigma_\epsilon^2)$$

The more experienced returns you have (higher s), the more precise is your memory signal. This is because there are two components: the noise coming from σ_ϵ^2 doesn't change with time, but the precision of the "informative" component, namely the experienced returns, increases with the number of participation periods.

To capture the second feature of human behaviour (incomplete information about ability and learning about ability using experiences), we assume that individuals do not know their true ability at the point of entry and that they update their beliefs of being a low ability investor using Bayesian updating. Let b_{it} denote the end-of-period belief of individual i who participated at time t that he is a low ability investor. By Bayes rule (derivation in Section C.2):

$$b_{it} = \frac{a_l \exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_l}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right]}{a_l \exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_l}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right] + (1-a_l)\exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_h}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right]} \tag{3}$$

Given a belief of being of low ability type, we can compute the expected return for individual i from investing in period $t+1$ as:

$$E_{i,t+1} = b_{it}\mu_l + (1-b_{it})\mu_h$$

The standard unconstrained model (Merton (1969)) says that any risk-averse investor with twice differentiable concave utility function should participate in risky financial assets in time $t+1$ as long as $E_{i,t+1} > r_s$, i.e. the expected risk premium is positive. This can be summarised in the following condition - you participate in $t+1$ iff:

$$b_{it} < \frac{\mu_h - r_s}{\mu_h - \mu_l}$$

Assume that the prior belief of being low type is given by the share of new entrants that are of low ability in the population, $b_{i0} = a_l$. To get everyone to participate initially, we need to assume:

$$a_l < \frac{\mu_h - r_s}{\mu_h - \mu_l}$$

If this were not true, then no-one would ever participate. Therefore, to summarise, individuals participate as long as:

$$b_{it} < \frac{\mu_h - r_s}{\mu_h - \mu_l}$$

where

$$b_{it} = \frac{a_l \exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_l}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right]}{a_l \exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_l}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right] + (1-a_l)\exp\left[-\frac{1}{2}\left(\frac{m_{it}-\mu_h}{\sqrt{\frac{\sigma^2}{s}+\sigma_\epsilon^2}}\right)^2\right]}$$

5.4 Calibration

To match the number of years for which we have Norwegian data, we take $T = 25$.[46] We also set the number of new entrants in each period, N_t, equal to the number of entrants in the corresponding year in the Norwegian data.[47] We then have six remaining parameters in the model: r_s, μ_h, μ_l, σ, σ_ϵ and a_l. The risk-free rate (r_s) is calibrated to equal the average 3-month Treasury bill rate in Norway over the period 1994-2018 (3.4%).[48] To calibrate the remaining five parameters, we use two external moments and three internal moments. The two external moments are the mean (4.25%) and standard deviation (24.73%) of returns on risky financial assets computed in Fagereng et al. (2020). As our three internal moments, we take the mean spell length across non-censored spells (5.58 years), the standard deviation of spell lengths across non-censored spells (4.84 years) and the probability of re-entry within 4 years (33%). We use a method of moments calibration strategy, whereby we find the set of five parameters, Θ, that solves:

$$\hat{\Theta} = \underset{\Theta}{\arg\min} \sum_{j=1}^{5} \left(\frac{m_j(\Theta) - \hat{m}_j}{\hat{m}_j}\right)^2 \tag{4}$$

[46]We cannot observe entry in 1993 as it is the first year in our dataset, and so we use 25 rather than 26 periods in the model.

[47]Note that we here we set the number of new entrants in each period in the model equal to the number of entrants (both re-entrants and new entrants) in the Norwegian data. Ideally, we would want to distinguish the two types in all periods and just use the number of new entrants; however, this is only realistic at the end of the sample once we have followed these individuals for a sufficient number of years.

[48]This is calculated using the average of monthly data obtained from Thomson Reuters Eikon.

where $\hat{m}_j$ is the j-th empirical moment targeted in the calibration and $m_j(\Theta)$ is the simulated moment from the model generated by parameter values Θ. As such, we minimise the sum of squared percentage deviations of the simulated model moment from the corresponding target empirical moment. Table 4 shows that the model-generated moments fit the targeted empirical moments well.

TABLE 4: Simulated vs. target empirical moments

Moment	Data	Model
External moments		
Average return (Fagereng et al. (2020))	0.043	0.042
Std. dev of returns (Fagereng et al. (2020))	0.247	0.247
Internal moments		
Average spell length (uncensored)	5.58	5.61
Std. dev of spell lengths (uncensored)	4.84	4.80
Re-entry within 4 years	0.33	0.33

Note: this table shows the performance of the method of moments calibration approach in Equation 4. The first column lists the five target moments. The first two moments are external moments: the mean and standard deviation of returns to risky financial assets reported in Table 3 of Fagereng et al. (2020). The remaining three moments are internal moments: the mean and standard deviation of spell lengths across uncensored spells and the proportion of exiters re-entering within 4 years. The second column gives the target value from the data and the third column gives the model-generated moment based on the optimal parameter values given in Table 5.

Table 5 summarises the parameter values that will be used in the simulations. The high ability individuals have an average excess return of 3.2%, while the low ability individuals have an average excess return of -2%. There are approximately equal proportions of high and low ability individuals across sets of new entrants with 50.6% of new entrants being of low ability. The standard deviation of returns is 24.59%, which can generate a large dispersion in experienced returns and means substantial overlap between the return distributions of the high and low ability types. The standard deviation on the memory noise is quite small at 1%.

5.5 Model simulations

Using the parameter values given in Table 5, we simulate the model and see whether it is able to generate the empirical facts described in Section 3. Figure 13 gives the distribution of spell lengths in this simulated sample. The red dots give the actual proportions from the Norwegian data. We are able to obtain the patterns observed in Figure 1. Quantitatively, we obtain slightly fewer short spells and too many right-censored observations in the model

TABLE 5: Parameter values

Parameter	Description	Method	Value
r_s	Risk-free return	External	0.034
μ_h	Mean return of high ability type	Internal	0.0656
μ_l	Mean return of low ability type	Internal	0.0142
a_l	Share of low ability individuals	Internal	0.5055
σ	Standard deviation of returns	Internal	0.2459
σ_ϵ	Standard deviation of memory noise	Internal	0.0100

Note: this table shows the parameter values, both those that are externally calibrated and those that are internally fitted based on the method of moments procedure (Equation 4).

relative to the data; however, we still find that short spells occur for a non-negligible proportion of participants. It is worth noting that the mean and standard deviation of uncensored spells, which are used as internal target moments, do not force the simulated histogram to take this declining shape. Indeed, many different shapes are consistent with the targeted mean and standard deviation.

Why do short spells occur in the model? This is primarily through individuals drawing poor returns and inferring low ability from their realised returns. This experience of bad returns causes individuals to update upwards their belief that they are of low ability. If the experience is sufficiently bad, then the belief will be pushed above the participation threshold and they will exit. Right-censored observations occur in the model because over time, if you receive consistently good returns, then you will be very confident that you are of high ability and so your belief of being low type, b_{it}, will be close to zero. As such, these individuals require a very poor return to convince them otherwise and drive them out of the market, which is very unlikely. In principle, some exit can also occur due to the memory noise. An individual could have had decent returns, but a poor recollection of returns may mean that they exit. However, the standard deviation of the memory noise is quite low and thus this is likely to be a secondary driver for exit because this margin of exit is only likely to be relevant for those who have experienced returns that place them close to the participation threshold. Figure D.21 shows the distribution of spell lengths by ability group in the model. Insofar as ability is linked to financial literacy, we are able to generate the empirical fact that short spells are more prevalent among individuals with lower financial literacy.

Figure 14 plots the hazard rates for the high and low ability groups separately. As ability is the only source of heterogeneity in the model, controlling for ability means the plotted

FIGURE 13: Spell length distribution

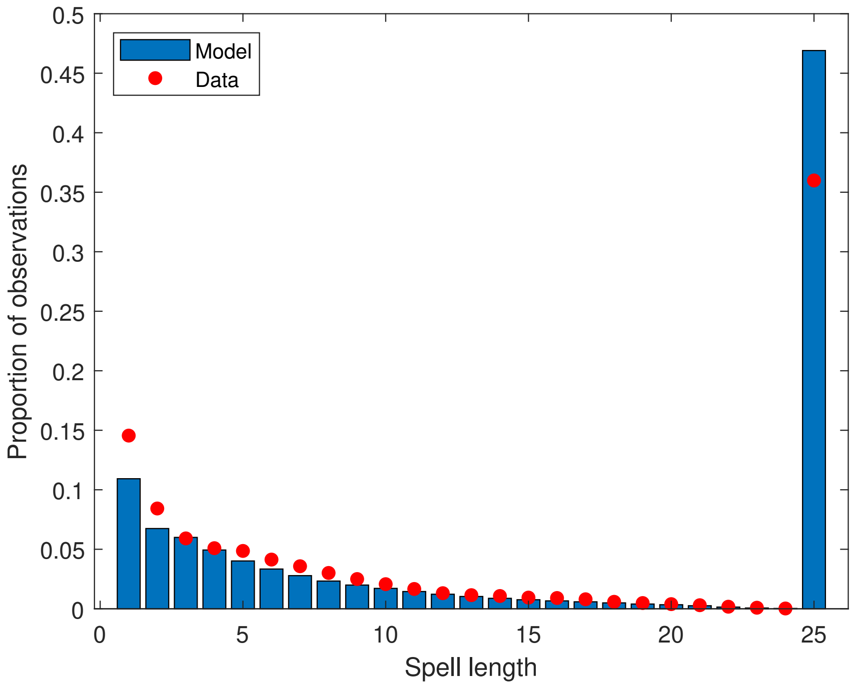

Note: this figure plots the distribution of spell lengths in the simulated model. The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data. The blue bars show the model simulation, while the red dots give the corresponding proportion in the Norwegian data. The bar at $T = 25$ is the proportion of right-censored spells.

hazard rates reflect true duration dependence. We observe very similar patterns to the estimated baseline hazard in Figure 3. In particular, for both ability groups the hazard function is downward sloping and generally convex. We see a sharp fall in the hazard rate by about 25% for both groups going from $d = 1$ to $d = 2$ compared to a 40% fall in Figure 3. The reason why the model is able to generate a downward-sloping hazard function is through the learning process. The hazard rate tells us the probability of leaving the market conditional on not having left until then. Suppose that you have been participating for 15 years: the fact that you have not yet left the market must mean that you have performed well so far, which means you should be reasonably confident that you are of high ability. It thus takes a very bad return realisation (or a very bad recollection noise) to drive you out of the market, which is a very low probability event and thus the hazard rate is low. In contrast, during the first few periods in the market, the only information available are the initial returns (or more precisely, recollections of the returns). As such, it does not take as bad a return to make you exit, meaning the hazard rate is higher at lower durations.

FIGURE 14: Hazard function by ability group

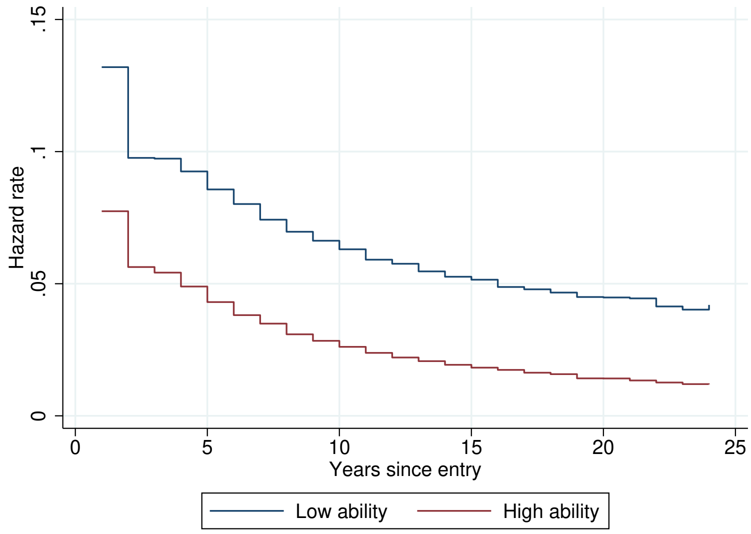

Note: this figure plots the hazard rates for exit from participation separately for the low and high ability groups in the simulated model. The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data. Hazard rates at duration d are computed as the proportion of individuals who are "at risk" of exit at duration d who do exit at this point, and is equivalent to the difference in cumulative hazard rates obtained from the Nelson-Aalen cumulative hazard estimator.

Moving onto the re-entry facts, Figure 15 plots the number of spells individuals have in the simulated model. As everyone participates in the first period, we compare this against the corresponding figure for only those who participate in the Norwegian data (Figure 4b). The model simulation fits the patterns in the Norwegian data very closely. How do multiple spells occur in the model? This is through the presence of noisy memory. Individuals do not remember precisely what their (average) experienced returns are - they recall their experiences with some noise. As such, re-entry occurs primarily amongst those individuals who have beliefs close to the participation threshold. These individuals are likely to have experienced moderately poor returns, but their fuzzy recollection may mean that at some point after exit, they think that they did sufficiently well to warrant their re-entry into the stock market. The model is also able to generate single spellers. Part of these will be right-censored participants; however, some will be individuals who did so badly that even with some imperfect recollection of their exact return, they will still conclude that they are really bad investors and stay out. Therefore, the model suggests that those who re-enter are less likely to have been burned in the market. To give some empirical support to this, we find that 30% of exiters who report only taxable gains re-enter within 4 years compared to 27% for exiters who report only losses. It is also supported by the behavioural/psychology literature which says that salient events are remembered well, but less salient events are less well remembered (Neligh (2021)). This idea would suggest that those who got burned should not forget and so they will stay out, while those who did moderately badly will get drawn back into the market.

Figure 16 shows the distribution of re-entry times. We almost exactly match the distribution in the data using the simulated model. The model is able to generate this behaviour because re-entry would be primarily driven by individuals with moderate returns who have beliefs close to the threshold. As such, for these individuals it takes just one positive recollection to drive them over the threshold and this is fairly likely within a few years following exit. In contrast, those who did terribly in their previous spell would strongly believe that they are of low ability. Even with noise, it is very hard to drive them past the threshold to participate because, for example, whether you recall a return of -8% or -10%, both are very poor returns that would deter re-entry. As such, re-entry at longer horizons is not likely.

Figure 17 plots the hazard function for re-entry following exit for the two ability groups separately. We are able to endogenously generate a downward-sloping re-entry hazard function akin to the baseline hazard estimated in Figure 7. The hazard rates drop by almost one-half for both groups from $d = 1$ to $d = 2$ compared to a 60% drop in Figure 7. The model can produce a downward-sloping re-entry hazard function because there is selection of who is likely to re-enter and how long it will take them. As discussed previously, re-entry will be more common for individuals who did moderately badly and are thus close to the threshold. Because of noisy recollection, they will soon after exit get a memory signal that induces them to re-enter. However, those who remain after say 10 years are likely exiters who did poorly and thus are so certain they are of low type that even with some noise in their recollection, they would never reach the participation threshold of beliefs. The hazard rates are greater for high ability compared to low ability individuals in the first few years after exit, which is in accordance with the findings in Fact 3 that re-entry is more common for individuals with characteristics associated with higher financial literacy. This arises because low

FIGURE 15: Number of spells

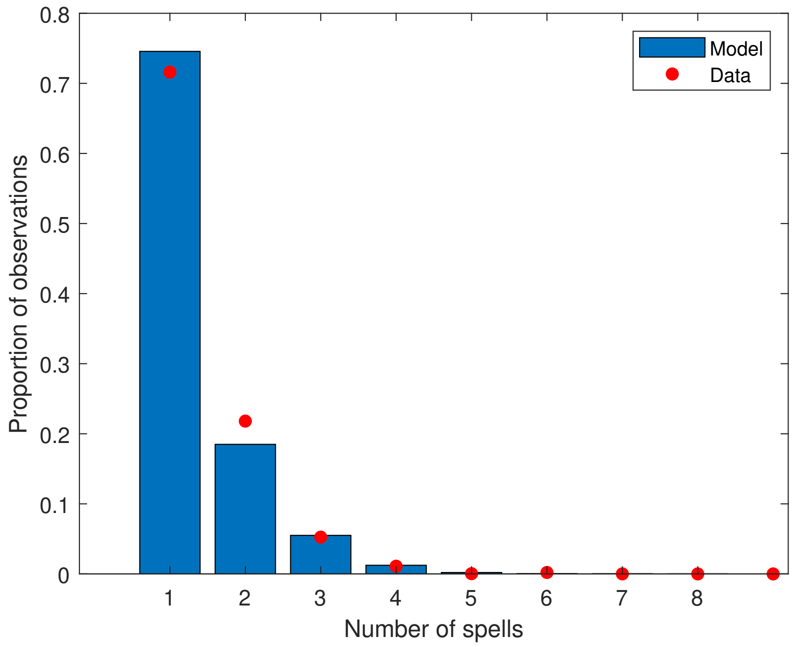

Note: this figure plots the proportion of individuals in the simulated model having s spells for different s. The red dots give the corresponding proportions from the Norwegian data (only looking at individuals who have had at least one spell). The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data.

ability individuals draw from a returns distribution with a lower mean. As such, they are less likely to have beliefs at the point of exit fairly close to the threshold because they are more likely to draw very poor returns and thus strongly believe that they are of low ability. Beyond this point, the hazard rates are effectively equivalent.

6 Conclusion

While there has been a large literature focused on understanding why, in contrast to the predictions of basic portfolio choice models, the participation rate at each static point in time is not 100%, much less is known about the dynamics of stock market participation by retail investors. How long do individuals stay in the stock market for? Is the probability of exit a function of time since entry? Do individuals re-enter after exit, and if so, when? To be able to study such questions empirically, panel data on individual wealth holdings that spans a sufficiently long time dimension is required. This paper documents five new em-

FIGURE 16: Time until re-entry

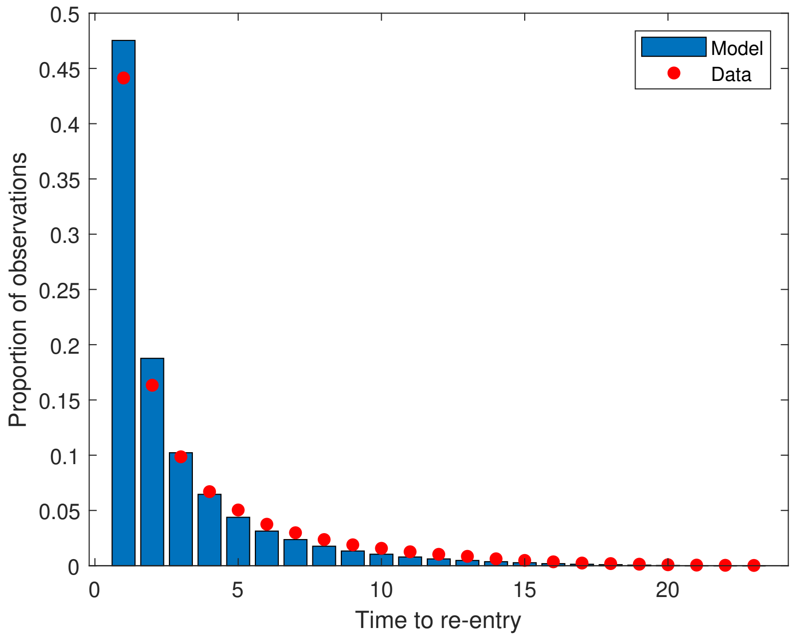

Note: this figure plots the distribution of re-entry times in the simulated model. The red dots give the corresponding proportions from the Norwegian data. The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data.

pirical facts on the dynamics of stock market participation using Norwegian administrative data that provides reliable and accurate information on wealth holdings for each member of the population: first, just under a quarter of all stock market spells end within just 2 years. These short spells occur across all population subgroups, but are more prevalent for groups with characteristics associated with lower financial literacy (no college education and low income/wealth). Second, we find evidence of negative duration dependence in exit from participation, which suggests that the longer you have been participating for, the lower is the probability of exiting. Third, $\approx$ 30% of exiters re-enter within 4 years of exit with re-entry being more common for those with a college degree or with high income/wealth. Fourth, conditional on occurring, re-entry typically happens very soon after exit, often just 1 year later. Fifth, there is negative duration dependence in re-entry following exit, which means that the longer you have been out of the market, the lower is the probability of re-entering. Overall, the broad message from the empirical facts is that short, multiple spells in the stock market are common. We then show that while standard classes of participation models are

FIGURE 17: Re-entry hazard function by ability group

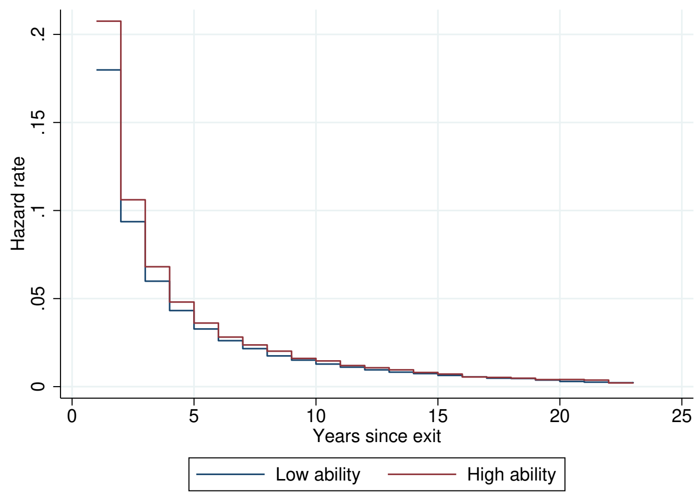

Note: this figure plots the hazard rates for re-entry following exit separately for the low and high ability groups in the simulated model. The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data. Hazard rates at duration d are computed as the proportion of individuals who are "at risk" of re-entry at duration d who do re-enter at this point, and is equivalent to the difference in cumulative hazard rates obtained from the Nelson-Aalen cumulative hazard estimator.

unlikely to be consistent with the dynamics observed in the data, our empirical findings are consistent with a model where investors use experienced returns, which are recalled with noise, to learn about their ex-ante heterogeneous ability.

There are various avenues for future research: first, undertaking similar analysis using data from other countries would be useful to establish whether such behaviours are present elsewhere. Second, the Norwegian data lacks information on the specific mutual funds held, which makes it difficult to obtain a clear idea of the nature of the individual portfolios given that most participants hold mutual funds rather than directly-held stocks. It would be interesting to understand how the nature of portfolios changes across spells. For example, do individuals who re-enter invest in the same firms or sectors as in their previous spell? Third, while the neuroscience and psychology literatures have established that memory is imperfect, there is little work testing imperfect memory in the context of financial markets. Indeed, our model applies this notion of imperfect memory and proposes it as a possible justification for re-entry. Further work trying to see how well (former) participants recall their past return experiences would thus help to establish whether noisy memory is a feature of investor behaviour. Fourth, the paper has mainly focused on the extensive margin of participation. Understanding how the intensive margin changes over time and across spells could also be insightful. Last, it would be interesting to understand the aggregate implications of short spells and quick re-entry. Do these transitions in and out of participation have effects on asset prices or wealth inequality dynamics?

References

ABEL, A. B., J. C. EBERLY, AND S. PANAGEAS (2007): "Optimal Inattention to the Stock Market," *American Economic Review*, 97(2), 244–249.

——— (2013): "Optimal inattention to the stock market with information costs and transaction costs," *Econometrica*, 81(4), 1455–1481.

AGNEW, J., P. BALDUZZI, AND A. SUNDÉN (2003): "Portfolio Choice and Trading in a Large 401(k) Plan," *American Economic Review*, 93(1), 193–215.

ALESINA, A., AND N. FUCHS-SCHÜNDELN (2007): "Good-bye Lenin (or Not?): The Effect of Communism on People's Preferences," *American Economic Review*, 97(4), 1507–1528.

ALVAREZ, F. E., K. BOROVIČKOVÁ, AND R. SHIMER (2021): "Consistent Evidence on Duration Dependence of Price Changes," Working Paper 29112, National Bureau of Economic Research.

AMERIKS, J., AND S. P. ZELDES (2004): "How do household portfolio shares vary with age," Discussion paper, Columbia University.

ANAGOL, S., V. BALASUBRAMANIAM, AND T. RAMADORAI (2021): "Learning from noise: Evidence from India's IPO lotteries," *Journal of Financial Economics*, 140(3), 965–986.

ANG, A., G. BEKAERT, AND J. LIU (2005): "Why stocks may disappoint," *Journal of Financial Economics*, 76(3), 471–508.

AZEREDO DA SILVEIRA, R., Y. SUNG, AND M. WOODFORD (2020): "Optimally imprecise memory and biased forecasts," Discussion paper, National Bureau of Economic Research.

BACH, L., L. E. CALVET, AND P. SODINI (2020): "Rich Pickings? Risk, Return, and Skill in Household Wealth," *American Economic Review*, 110(9), 2703–47.

BARBER, B. M., Y.-T. LEE, Y.-J. LIU, AND T. ODEAN (2008): "Just How Much Do Individual Investors Lose by Trading?," *The Review of Financial Studies*, 22(2), 609–632.

BARBER, B. M., Y.-T. LEE, Y.-J. LIU, T. ODEAN, AND K. ZHANG (2019): "Learning, Fast or Slow," *The Review of Asset Pricing Studies*, 10(1), 61–93.

BARBER, B. M., AND T. ODEAN (2000): "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *The Journal of Finance*, 55(2), 773–806.

——— (2001): "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," *The Quarterly Journal of Economics*, 116(1), 261–292.

BARBERIS, N., R. GREENWOOD, L. JIN, AND A. SHLEIFER (2018): "Extrapolation and bubbles," *Journal of Financial Economics*, 129(2), 203–227.

BASTEN, C., A. FAGERENG, AND K. TELLE (2016): "Saving and Portfolio Allocation Before and After Job Loss," *Journal of Money, Credit and Banking*, 48(2-3), 293–324.

BEHRMAN, J. R., O. S. MITCHELL, C. K. SOO, AND D. BRAVO (2012): "How Financial Literacy Affects Household Wealth Accumulation," *American Economic Review*, 102(3), 300–304.

BIANCHI, M. (2018): "Financial Literacy and Portfolio Dynamics," *The Journal of Finance*, 73(2), 831–859.

BONAPARTE, Y., G. M. KORNIOTIS, AND A. KUMAR (2021): "Income Risk and Stock Market Entry/Exit Decisions," Discussion paper, CEPR Discussion Paper No. DP15370.

BRANDSAAS, E. E. (2021): "Household Stock Market Participation and Exit: The Role of Homeownership," Working paper.

BROCAS, I., AND J. D. CARRILLO (2016): "A neuroeconomic theory of memory retrieval," *Journal of Economic Behavior & Organization*, 130, 198–205.

BROWN, J. R., Z. IVKOVIĆ, P. A. SMITH, AND S. WEISBENNER (2008): "Neighbors Matter: Causal Community Effects and Stock Market Participation," *The Journal of Finance*, 63(3), 1509–1531.

BRUNNERMEIER, M. K., AND S. NAGEL (2008): "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-evidence on Individuals," *American Economic Review*, 98(3), 713–36.

CALVET, L. E., J. Y. CAMPBELL, AND P. SODINI (2007): "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 115(5), 707–747.

——— (2009a): "Fight or Flight? Portfolio Rebalancing by Individual Investors," *The Quarterly Journal of Economics*, 124(1), 301–348.

——— (2009b): "Measuring the Financial Sophistication of Households," *American Economic Review*, 99(2), 393–98.

CAMPBELL, J. Y. (2006): "Household Finance," *The Journal of Finance*, 61(4), 1553–1604.

CAO, H. H., T. WANG, AND H. H. ZHANG (2005): "Model Uncertainty, Limited Market Participation, and Asset Prices," *The Review of Financial Studies*, 18(4), 1219–1251.

CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23(2), 61–156.

CARROLL, C. D., B. RHEE, AND C. RHEE (1999): "Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants," *Economic Development and Cultural Change*, 48(1), 33–50.

CARROLL, C. D., B.-K. RHEE, AND C. RHEE (1994): "Are There Cultural Effects on Saving? Some Cross-Sectional Evidence*," *The Quarterly Journal of Economics*, 109(3), 685–699.

CHARNESS, G., AND D. LEVIN (2005): "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect," *American Economic Review*, 95(4), 1300–1309.

CHIANG, Y.-M., D. HIRSHLEIFER, Y. QIAN, AND A. E. SHERMAN (2011): "Do Investors Learn from Experience? Evidence from Frequent IPO Investors," *The Review of Financial Studies*, 24(5), 1560–1589.

CHRISTIANSEN, C., J. S. JOENSEN, AND J. RANGVID (2015): "Understanding the effects of marriage and divorce on financial investments: the role of background risk sharing," *Economic Inquiry*, 53(1), 431–447.

CHUANG, Y., AND L. SCHECHTER (2015): "Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results," *Journal of Development Economics*, 117, 151–170.

COCCO, J. F., F. J. GOMES, AND P. J. MAENHOUT (2005): “Consumption and Portfolio Choice over the Life Cycle,” *The Review of Financial Studies*, 18(2), 491–533.

COHN, A., J. ENGELMANN, E. FEHR, AND M. A. MARÉCHAL (2015): “Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals,” *American Economic Review*, 105(2), 860–85.

COOPER, J. (2019): “How does UK healthcare spending compare with other countries,” *Office for National Statistics*.

DAVIS, S. J., AND P. WILLEN (2013): “Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice,” *The Quarterly Journal of Finance*, 03(03n04), 1350011.

DEUFLHARD, F., D. GEORGARAKOS, AND R. INDERST (2018): “Financial Literacy and Savings Account Returns,” *Journal of the European Economic Association*, 17(1), 131–164.

DIMSON, E., P. MARSH, AND M. STAUNTON (2008): “The Worldwide Equity Premium: A Smaller Puzzle,” in *Handbook of the Equity Risk Premium*, ed. by R. Mehra, Handbooks in Finance, pp. 467–514. Elsevier.

DOHMEN, T., A. FALK, B. H. H. GOLSTEYN, D. HUFFMAN, AND U. SUNDE (2017): “Risk Attitudes Across The Life Course,” *The Economic Journal*, 127(605), F95–F116.

DOHMEN, T., H. LEHMANN, AND N. PIGNATTI (2016): “Time-varying individual risk attitudes over the Great Recession: A comparison of Germany and Ukraine,” *Journal of Comparative Economics*, 44(1), 182–200, Ukraine Escape from Post-Soviet Legacy.

DOMINITZ, J., AND C. F. MANSKI (2011): “Measuring and interpreting expectations of equity returns,” *Journal of Applied Econometrics*, 26(3), 352–370.

EBBINGHAUS, H. (1885): *Über das gedächtnis: untersuchungen zur experimentellen psychologie*. Duncker & Humblot.

EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4), 937–969.

FAGERENG, A., C. GOTTLIEB, AND L. GUISO (2017a): “Asset Market Participation and Portfolio Choice over the Life-Cycle,” *The Journal of Finance*, 72(2), 705–750.

FAGERENG, A., L. GUISO, D. MALACRINO, AND L. PISTAFERRI (2020): “Heterogeneity and Persistence in Returns to Wealth,” *Econometrica*, 88(1), 115–170.

FAGERENG, A., L. GUISO, AND L. PISTAFERRI (2017b): “Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk,” *The Review of Economic Studies*, 85(1), 437–474.

FAGERENG, A., M. B. HOLM, B. MOLL, AND G. NATVIK (2019): “Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains,” Working Paper 26588, National Bureau of Economic Research.

FORGAS, J. (1995): “Mood and judgment: the affect infusion model (AIM),” *Psychological bulletin*, 117(1), 39—66.

FREDRIKSEN, D., AND E. HALVORSEN (2019): “Beregninger av pensjonsformue,” Discussion paper, Statistics Norway.

FREDRIKSEN, D., AND N. M. STØLEN (2018): “Reform av offentlig tjenestepensjon,” Discussion paper, Statistics Norway.

FUCHS-SCHÜNDELN, N., P. MASELLA, AND H. PAULE, PALUDKIEWICZ (2020): “Cultural Determinants of Household Saving Behavior,” *Journal of Money, Credit and Banking*, 52(5), 1035–1070.

GABAIX, X., J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2016): “The Dynamics of Inequality,” *Econometrica*, 84(6), 2071–2111.

GAUDECKER, H.-M. V. (2015): “How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?,” *The Journal of Finance*, 70(2), 489–507.

GOMES, F. (2020): “Portfolio Choice Over the Life Cycle: A Survey,” *Annual Review of Financial Economics*, 12(1), 277–304.

GOMES, F., M. HALIASSOS, AND T. RAMADORAI (2021): “Household Finance,” *Journal of Economic Literature*, 59(3), 919–1000.

GOMES, F., AND A. MICHAELIDES (2005): “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *The Journal of Finance*, 60(2), 869–904.

GOMES, F., AND O. SMIRNOVA (2021): “Stock Market Participation and Portfolio Shares Over the Life-Cycle,” *Available at SSRN 3808350*.

GREENWOOD, R., AND S. NAGEL (2009): “Inexperienced investors and bubbles,” *Journal of Financial Economics*, 93(2), 239–258.

GRINBLATT, M., M. KELOHARJU, AND J. LINNAINMAA (2011): “IQ and Stock Market Participation,” *The Journal of Finance*, 66(6), 2121–2164.

GUIN, B. (2017): “Culture and household saving,” Working Paper Series 2069, European Central Bank.

GUISO, L., P. SAPIENZA, AND L. ZINGALES (2003): “People’s opium? Religion and economic attitudes,” *Journal of Monetary Economics*, 50(1), 225–282.

——— (2006): “Does Culture Affect Economic Outcomes?,” *Journal of Economic Perspectives*, 20(2), 23–48.

——— (2018): “Time varying risk aversion,” *Journal of Financial Economics*, 128(3), 403–421.

GUISO, L., AND P. SODINI (2013): “Chapter 21 - Household Finance: An Emerging Field,” vol. 2 of *Handbook of the Economics of Finance*, pp. 1397–1532. Elsevier.

HALIASSOS, M., AND C. C. BERTAUT (1995): “Why do so Few Hold Stocks?,” *The Economic Journal*, 105(432), 1110–1129.

HALIASSOS, M., T. JANSSON, AND Y. KARABULUT (2017): “Incompatible European Partners? Cultural Predispositions and Household Financial Behavior,” *Management Science*, 63(11), 3780–3808.

HANSEN, L. P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50(4), 1029–1054.

HECKMAN, J., AND B. SINGER (1984): "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 52(2), 271–320.

HONG, H., J. D. KUBIK, AND J. C. STEIN (2004): "Social Interaction and Stock-Market Participation," *The Journal of Finance*, 59(1), 137–163.

HONORÉ, B. E. (1993): "Identification Results for Duration Models with Multiple Spells," *Review of Economic Studies*, 60(1), 241–46.

HUDOMIET, P., G. KÉZDI, AND R. J. WILLIS (2011): "Stock market crash and expectations of American households," *Journal of Applied Econometrics*, 26(3), 393–415.

HURD, M., M. VAN ROOIJ, AND J. WINTER (2011): "Stock market expectations of Dutch households," *Journal of Applied Econometrics*, 26(3), 416–436.

HURST, E., M. C. LUOH, F. P. STAFFORD, AND W. G. GALE (1998): "The Wealth Dynamics of American Families, 1984-94," *Brookings Papers on Economic Activity*, 1998(1), 267–337.

JAGANNATHAN, R., AND N. KOCHERLAKOTA (1996): "Why Should Older People Invest Less in Stocks Than Younger People?," *Quarterly Review*.

KAHNEMAN, D., AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–291.

KANDASAMY, N., B. HARDY, L. PAGE, M. SCHAFFNER, J. GRAGGABER, A. S. POWLSON, P. C. FLETCHER, M. GURNELL, AND J. COATES (2014): "Cortisol shifts financial risk preferences," *Proceedings of the National Academy of Sciences*, 111(9), 3608–3613.

KARLSSON, N., G. LOEWENSTEIN, AND D. SEPPI (2009): "The ostrich effect: Selective attention to information," *Journal of Risk and uncertainty*, 38(2), 95–115.

KAUSTIA, M., AND S. KNÜPFER (2008): "Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions," *The Journal of Finance*, 63(6), 2679–2702.

KAUSTIA, M., AND S. KNÜPFER (2012): "Peer performance and stock market entry," *Journal of Financial Economics*, 104(2), 321–338, Special Issue on Investor Sentiment.

KIEFER, N. M. (1988): "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 26(2), 646–679.

KLOVLAND, J. T. (2004): "Bond markets and bond yields in Norway 1820–2003," in *Historical Monetary Statistics for Norway 1819-2003*, ed. by Øyvind Eitrheim, J. T. Klovland, and J. F. Qvigstad, no. 35. Norges Bank.

KNÜPFER, S., E. RANTAPUSKA, AND M. SARVIMÄKI (2017): "Formative Experiences and Portfolio Choice: Evidence from the Finnish Great Depression," *The Journal of Finance*, 72(1), 133–166.

KROFT, K., F. LANGE, M. J. NOTOWIDIGDO, AND L. F. KATZ (2016): "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation," *Journal of Labor Economics*, 34(S1), S7–S54.

KUCHLER, T., AND B. ZAFAR (2019): "Personal Experiences and Expectations about Aggregate Outcomes," *The Journal of Finance*, 74(5), 2491–2542.

LANCASTER, T. (1979): "Econometric Methods for the Duration of Unemployment," *Econometrica*, 47(4), 939–956.

LEVIN, I. P., J. A. WELLER, A. A. PEDERSON, AND L. A. HARSHMAN (2007): "Age-related differences in adaptive decision making: Sensitivity to expected value in risky choice," *Judgment and Decision Making*, 2(4), 225–233.

LINNAINMAA, J. T. (2011): "Why Do (Some) Households Trade So Much?," *The Review of Financial Studies*, 24(5), 1630–1666.

LUSARDI, A., AND O. S. MITCHELL (2011): "Financial Literacy and Planning: Implications for Retirement Well-being," in *Financial Literacy: Implications for Retirement Security and the Financial Marketplace*. Oxford University Press, Oxford.

MAHANI, R., AND D. BERNHARDT (2007): "Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity," *The Journal of Finance*, 62(3), 1313–1340.

MALMENDIER, U., AND S. NAGEL (2011): "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?," *The Quarterly Journal of Economics*, 126(1), 373–416.

——— (2015): " Learning from Inflation Experiences," *The Quarterly Journal of Economics*, 131(1), 53–87.

MALMENDIER, U., AND G. TATE (2005): "CEO Overconfidence and Corporate Investment," *The Journal of Finance*, 60(6), 2661–2700.

MANKIW, N., AND S. P. ZELDES (1991): "The consumption of stockholders and nonstockholders," *Journal of Financial Economics*, 29(1), 97–112.

MERTON, R. C. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, 51(3), 247–257.

MERTON, R. C. (1971): "Optimum consumption and portfolio rules in a continuous-time model," *Journal of Economic Theory*, 3(4), 373–413.

MEYER, B. D. (1990): "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58(4), 757–782.

MUELLER, A. I., J. SPINNEWIJN, AND G. TOPA (2021): "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias," *American Economic Review*, 111(1), 324–63.

NAKAMURA, E., AND J. STEINSSON (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

NELIGH, N. (2021): "Rational Memory with Decay," Working paper.

NICOLOSI, G., L. PENG, AND N. ZHU (2009): "Do individual investors learn from their trading experience?," *Journal of Financial Markets*, 12(2), 317–336.

ODEAN, T. (1998): "Are Investors Reluctant to Realize Their Losses?," *The Journal of Finance*, 53(5), 1775–1798.

ØDEGAARD, B. A. (2007): "Asset Pricing at the Oslo Stock Exchange: A Source Book," *Available at SSRN 969390.*

OECD (2009): "OECD Private Pensions Outlook 2008," Discussion paper, OECD.

OREOPOULOS, P., T. VON WACHTER, AND A. HEISZ (2012): "The Short- and Long-Term Career Effects of Graduating in a Recession," *American Economic Journal: Applied Economics*, 4(1), 1–29.

OSILI, U. O., AND A. PAULSON (2008): "What Can We Learn about Financial Access from U.S. Immigrants? The Role of Country of Origin Institutions and Immigrant Beliefs," *The World Bank Economic Review*, 22(3), 431–455.

PAGEL, M. (2018): "A news-utility theory for inattention and delegation in portfolio choice," *Econometrica*, 86(2), 491–522.

PAULSEN, D. J., M. L. PLATT, S. A. HUETTEL, AND E. M. BRANNON (2012): "From risk-seeking to risk-averse: the development of economic risk preference from childhood to adulthood," *Frontiers in psychology*, 3, 313.

POLDRACK, R. A., AND K. FOERDE (2008): "Category learning and the memory systems debate," *Neuroscience Biobehavioral Reviews*, 32(2), 197–205, The Cognitive Neuroscience of Category Learning.

POTERBA, J. M., AND A. A. SAMWICK (1997): "Household Portfolio Allocation Over the Life Cycle," Working Paper 6185, National Bureau of Economic Research.

SAHM, C. R. (2012): "How Much Does Risk Tolerance Change?," *The Quarterly Journal of Finance*, 2(4).

SAMUELSON, P. A. (1969): "Lifetime Portfolio Selection By Dynamic Stochastic Programming," *The Review of Economics and Statistics*, 51(3), 239–246.

SCHILDBERG-HÖRISCH, H. (2018): "Are Risk Preferences Stable?," *Journal of Economic Perspectives*, 32(2), 135–54.

SCHURER, S. (2015): "Lifecycle patterns in the socioeconomic gradient of risk preferences," *Journal of Economic Behavior & Organization*, 119, 482–495.

SEGAL, U., AND A. SPIVAK (1990): "First order versus second order risk aversion," *Journal of Economic Theory*, 51(1), 111–125.

SERU, A., T. SHUMWAY, AND N. STOFFMAN (2010): "Learning by Trading," *The Review of Financial Studies*, 23(2), 705–739.

SHILLER, R. J., S. FISCHER, AND B. M. FRIEDMAN (1984): "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, 1984(2), 457–510.

SIMS, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50(3), 665–690, Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

STIGLER, G. J., AND G. S. BECKER (1977): "De Gustibus Non Est Disputandum," *The American Economic Review*, 67(2), 76–90.

TAUCHEN, G., AND R. HUSSEY (1991): "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59(2), 371–396.

VISSING-JØRGENSEN, A. (2002): "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," Working Paper 8884, National Bureau of Economic Research.

XAVIER, I. (2021): "Wealth inequality in the US: the role of heterogeneous returns," *Available at SSRN 3915439*.

Appendix

A Variable construction

Here we describe the steps undertaken to translate the tax records into consistent measures of wealth by broad asset class. TR x.y will denote item x.y in the tax records based on 2018 item codings by the Norwegian Tax Administration (Skatteetaten). Note that while tax values are reported in the raw data, we translate these values into market values for our analysis. For financial wealth, we create the following subclasses:

- Cash and deposits are computed as the sum of deposits in Norwegian banks (TR 4.1.1), cash (TR 4.1.3), deposits in foreign banks (TR 4.1.9) and (from 2017 onwards) cash holdings in share savings accounts (TR 4.1.8.6).
- Directly-held listed stocks are given by the value of listed Norwegian shares and equity certificates, bonds, etc. in the Norwegian Central Securities Depository (TR 4.1.7).
- Directly-held unlisted stocks are given by capital in unlisted shares, share savings accounts and securities not listed in the Norwegian Central Securities Depository (TR 4.1.8).
- Stock mutual fund holdings are given by the value of the share component in holdings of securities funds (TR 4.1.4) plus (from 2017 onwards) equity holdings in share savings accounts (TR 4.1.8.5).
- Money market/bond funds are given by the value of the interest component in holdings of securities funds (TR 4.1.5).
- Financial wealth held abroad is given by other taxable capital abroad such as foreign shares, outstanding claims, bonds and endowment insurance (TR 4.6.2).
- Other financial assets are the sum of outstanding receivables in Norway (TR 4.1.6), share of capital in housing cooperatives or jointly owned property (TR 4.5.3), own pension insurance and life insurance (TR 4.5.1 + TR 4.5.2) and other taxable capital such as cryptocurrency (TR 4.5.4).

For real wealth, we decompose into:

- Housing wealth is the sum of housing owned through housing cooperatives (TR 4.3.2.2) and self-owned property (TR 4.3.2.1 + TR 4.3.2.3).
- Other real wealth is the sum of boats (TR 4.2.4), cars (TR 4.2.5), caravans (TR 4.2.6), holiday homes (TR 4.3.3.1 + TR 4.3.2.3), other real estate (TR 4.3.4 + TR 4.3.5 + TR 4.3.2.3), home contents and movable property (TR 4.2.3), fixtures & other business assets (TR 4.4.1 + TR 4.4.2 + TR 4.4.3 + TR 4.4.4) and real wealth abroad (TR 4.6.1 + TR 4.3.6.1).

We then treat an individual as a participant in the stock market if any of directly-held listed stock holdings, stock mutual fund holdings or financial wealth held abroad are strictly positive.

B Details on the life-cycle model of Section 5.1

B.1 Model setup

B.1.1 Labour income process

Individuals have a finite horizon T that can be split into two periods: working age ($t \leq t_r$) and retirement ($t > t_r$). During working age, individuals receive a labour income subject to undiversifiable shocks, but in retirement they receive a fixed pension.

In working life, your labour income depends on a deterministic function of age $f(t)$ that will be calibrated to capture the hump-shaped nature of earnings during working life, as well as a transitory component u_{it} and a persistent component p_{it} modelled as a random walk. t_b reflects the age you are born at in the model.

$$\ln(Y_{it}) = f(t) + p_{it} + u_{it} \quad \text{for } t \in \{t_b, ..., t_r\},\ u_{it} \sim N(0, \sigma_u^2)$$

$$p_{it} = p_{i,t-1} + z_{it}\ ,\ z_{it} \sim N(0, \sigma_z^2)$$

The current level of permanent income Y^p_{it} is defined as:

$$Y^p_{it} \equiv \exp(p_t)\exp(f(t))$$

During retirement, agents receive an income that is a fraction ϕ_{ret} of their permanent income in the last year of working life. Note this means that once you reach retirement, you face no uncertainty over your labour income.

$$\ln(Y_{it}) = \ln(\phi_{ret}) + \ln(Y^p_{it_r}) = \ln(\phi_{ret}) + f(t_r) + p_{it_r} \quad \text{for}\, t \in \{t_r+1, ..., T\} \tag{5}$$

B.1.2 Preferences

Households face an Epstein-Zin utility function (Epstein and Zin (1989)) over consumption:

$$U_{it} = \left[(1-\beta)C_{it}^{1-\frac{1}{\psi}} + \beta E_t(\pi_t U_{i,t+1}^{1-\gamma})^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\psi}}}$$

ψ is the elasticity of intertemporal substitution, β is the subjective discount factor and γ is the coefficient of relative risk aversion. π_t is the conditional survival probability, i.e. the probability of surviving to age $t+1$ conditional on being alive at age t.

B.1.3 Financial markets and participation costs

Individuals can invest in a riskless bond with safe return R_f or a risky asset (stocks) with a stochastic return R_{it} that follows the process:

$$1 + R_{it} = 1 + R_f + \bar{R} + \epsilon_{it} \quad \text{where}\ \epsilon \sim N(0, \sigma^2_\epsilon)$$

where $\bar{R}$ denotes the average equity risk premium.

Unlike in the workhorse Cocco et al. (2005) model, we add stock market participation costs. We follow Gomes and Smirnova (2021) and consider two costs: an entry cost (F^0) which has to be paid at the start of any new participation spell, and a per-period cost (F^1) paid in any period where you choose a non-zero quantity of stocks. We assume that these costs are proportional to the level of permanent income. This is a common assumption made in the literature (e.g. Gomes and Michaelides (2005)) because it simplifies the solving of the model. In particular, we are able to exploit scale invariance of the problem to scale the problem by permanent income and thus remove permanent income as a state variable. However, it can also be motivated by the view that participation costs reflect the opportunity cost of time, which is higher for those with higher permanent income.[49]

B.1.4 Optimisation problem

Let X_{it} denote your cash on hand at the start of age t to use for consumption or saving. Your choice variables are consumption C_{it} and the risky asset share α_{it}. Given choices, cash on hand follows the following process:

$$X_{i,t+1} = \big(\alpha_{it}(1+R_{i,t+1})+(1-\alpha_{it})(1+R_f)\big)\big(X_{it}-C_{it}-F^0 Y^p_{it}\mathbb{1}(S_{i,t-1}=0)-F^1 Y^p_{it}\mathbb{1}(S_{it}>0)\big)+Y_{i,t+1} \tag{6}$$

There are also borrowing ($\alpha_{it} \leq 1$) and no short selling ($\alpha_{it} \geq 0$) constraints, as well as a standard non-negativity constraint on consumption $C_{it} \geq 0$.

The individual's decision problem is then

$$\max_{\{C_{it},\alpha_{it}\}_{t=t_b}^{T}} E\sum_{t=t_b}^{T} \beta^{t-1}\Big(\prod_{j=t_b}^{t} \pi_j\Big)U_{it} \tag{7}$$

subject to the above constraints.

An advantage of the problem setup is that the value function is homogeneous with respect to current permanent labour income. This allows us to normalise by permanent labour income, meaning we can remove permanent income as a state variable when solving the model.[50] The state variables are then just age t and cash on hand X_{it}. We can write the Bellman equation as:

$$V_{it}(X_{it}) = \max_{C_{it}\geq 0, 0\leq\alpha_{it}\leq 1} U(C_{it}) + \beta\pi_{t+1}E_t V_{i,t+1}(X_{i,t+1}) \tag{8}$$

where

$$X_{i,t+1} = \big(\alpha_{it}(1+R_{i,t+1})+(1-\alpha_{it})(1+R_f)\big)\big(X_{it}-C_{it}-F^0 Y^p_{it}\mathbb{1}(S_{i,t-1}=0)-F^1 Y^p_{it}\mathbb{1}(S_{it}>0)\big)+Y_{i,t+1}$$

[49] Fagereng et al. (2017a) consider a nominal \$300 participation cost that does not vary across individuals. This would require permanent income to be a state variable, which complicates the solving of the model.

[50] See Carroll (1992) for an explanation and derivation of this.

This problem is solved using backward induction by noting that in the final period, you consume all your cash on hand ($C_{iT} = X_{iT}$) given that there is no bequest motive in this setup. The state space for cash on hand is discretised, and the return and labour income shocks are discretised using the method of Tauchen and Hussey (1991).

B.1.5 Calibration

Table 6 gives the parameter values used to calibrate the model. The values chosen are drawn from other studies and are based on Norwegian data. For the participation costs, we use the values of the "medium financial literacy" group in Gomes and Smirnova (2021). The entry cost F^0 is set as 3% of annual permanent income, while the per-period participation cost is 0.25%. The remaining parameters are set as follows: the subjective discount factor β is 0.96 and the coefficient of relative risk aversion γ is set at 5. We set the elasticity of intertemporal substitution equal to the inverse of the coefficient of relative risk aversion, so we effectively return to a power utility setting ($\gamma = \frac{1}{\psi}$). We take $t_b = 25$ and $T = 100$. We assume no correlations between the income shocks and stock returns.

TABLE 6: Calibrated parameter values

Parameter	Description	Value	Source
t_r	Retirement age	67	Norwegian law (Fagereng et al. (2017a))
F^0	Entry cost	3%	Medium group in Gomes and Smirnova (2021)
F^1	Per-period cost	0.25%	Medium group in Gomes and Smirnova (2021)
$f(t)$	Deterministic wage profile	-	Polynomial estimated in Fagereng et al. (2017a)
ϕ_{ret}	Replacement ratio	0.842	Fagereng et al. (2017a)
σ_z	Std. dev of permanent shock	0.110	Fagereng et al. (2017a)
σ_u	Std. dev of temporary shock	0.152	Fagereng et al. (2017a)
R_f	Risk-free return	0.0143	Klovland (2004)
$\bar{R}$	Average risk premium	0.0314	Dimson et al. (2008)
σ_ϵ	Std. dev of stock return	0.238	Ødegaard (2007)
π_t	Cond'l survival probabilities	-	SSB Life Tables 2010

Note: this table shows the calibrated parameter values used in the simulation of the life-cycle model described in Sections 5.1.

C Additional details

C.1 Further details on Alvarez et al. (2021) GMM estimator

The Alvarez et al. (2021) GMM estimator is based on the following environment: there is a proportional hazards data generating process for durations $d \in \{\underline{D}, ..., \bar{D}\}$ where $h_i(d) = \theta_i b_d$. Individual i experiences K^i spells, for which the measured duration of spells is $\zeta^i = \{\zeta^i_0, \zeta^i_1, ..., \zeta^i_{K^i}\}$. Note that measured duration is not necessarily equal to the true length of the spell because of censoring. Assume that the spells $\zeta = (\zeta_0, \zeta_1, ..., \zeta_K)$ are drawn from a proportional hazards model with a baseline hazard $\mathbf{b}_0$. Defining

$$f^{[b]}_{d_1,d_2}(\zeta;\mathbf{b}) \equiv \sum_{(j,k):1\leq j\leq k\leq K} (b_{d_2}\mathbb{1}_{\zeta_j=d_1,\zeta_k\geq d_2} - b_{d_1}\mathbb{1}_{\zeta_j=d_2,\zeta_k\geq d_1})$$

then $\mathbb{E}[f^{[b]}_{t_1,t_2}] = 0 \ \forall \ \underline{D} \leq d_1 < d_2 \leq \bar{D}$ if and only if $\mathbf{b} = \lambda \mathbf{b}_0$ for some $\lambda > 0$. This gives $\frac{\tilde{D}(\tilde{D}+1)}{2}$ moment conditions where $\tilde{D} \equiv \bar{D} - \underline{D}$. It is important to note that under this procedure, we recover the baseline hazards $\mathbf{b}$ up to a multiplicative constant, and so we normalise $b_1 = 1$. To estimate $\mathbf{b}_0$:

$$\hat{\mathbf{b}}_0 = \arg\min_{\mathbf{b}} \Big(\frac{1}{N}\sum_{i=1}^{N} f^{[b]}_{d_1,d_2}(\zeta^i;\mathbf{b})\Big)^T W \Big(\frac{1}{N}\sum_{i=1}^{N} f^{[b]}_{d_1,d_2}(\zeta^i;\mathbf{b})\Big)$$

where W is a positive definite weighting matrix. We use two-step feasible GMM à la Hansen (1982). In the first step, we use the identity matrix as the weighting matrix. In the second step, we take the estimates from the first step, $\mathbf{b}^{(1)}_0$, and use $\hat{W}(\hat{\mathbf{b}}_0)^{-1}$ as the weighting matrix in the second step where:[51]

$$\hat{W}(\hat{\mathbf{b}}_0) = \left(\frac{1}{N}\sum_{i=1}^{N} f^{[b]}_{d_1,d_2}(\zeta^i;\hat{\mathbf{b}}_0) f^{[b]}_{d_1,d_2}(\zeta^i;\hat{\mathbf{b}}_0)^T\right)^{-1}$$

[51] Hansen (1982) show that $\hat{W}(\hat{\mathbf{b}}_0)$ converges in probability to $\Omega \equiv \mathbb{E}[f^{[b]}_{d_1,d_2}(\zeta^i;\mathbf{b}_0) f^{[b]}_{d_1,d_2}(\zeta^i;\mathbf{b}_0)^T]$ and that $W = \Omega^{-1}$ is the most efficient weighting matrix.

C.2 Derivation of Equation 3

In a setting with discrete parameter values $\theta \in \Theta$ and a continuous observation $x \in X$, the Bayes rule formula is:

$$P_{\Theta|X}(\theta|x) = \frac{P_\Theta(\theta) \cdot f_{X|\Theta}(x|\theta)}{f_X(x)}$$

where $f_X(x) = \sum_\theta P_\Theta(\theta) \cdot f_{X|\Theta}(x|\theta)$. In our setting, we have:

$$b_{it} \equiv P(a_i = l|m_{it}) = \frac{P(a^i = l) \cdot f(m_{it}|a^i = l)}{f(m_{it})}$$

where

$$\begin{aligned}
P(a^i = l) &= a_l \\
f(m_{it}|a^i = l) &= \frac{1}{\sqrt{2\pi(\frac{\sigma^2}{s} + \sigma_\epsilon^2)}} \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_l}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big] \\
f(m_{it}) &= P(a^i = l) \cdot f(m_{it}|a^i = l) + P(a^i = h) \cdot f(m_{it}|a^i = h) \\
&= a_l \frac{1}{\sqrt{2\pi(\frac{\sigma^2}{s} + \sigma_\epsilon^2)}} \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_l}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big] + (1 - a_l) \frac{1}{\sqrt{2\pi(\frac{\sigma^2}{s} + \sigma_\epsilon^2)}} \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_h}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big]
\end{aligned}$$

Putting these terms all together, we get:

$$b_{it} = \frac{a_l \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_l}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big]}{a_l \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_l}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big] + (1 - a_l) \exp\Big[-\frac{1}{2}\Big(\frac{m_{it} - \mu_h}{\sqrt{\frac{\sigma^2}{s} + \sigma_\epsilon^2}}\Big)^2\Big]}$$

D Additional tables and figures

FIGURE D.1: Stock market participation rate over time

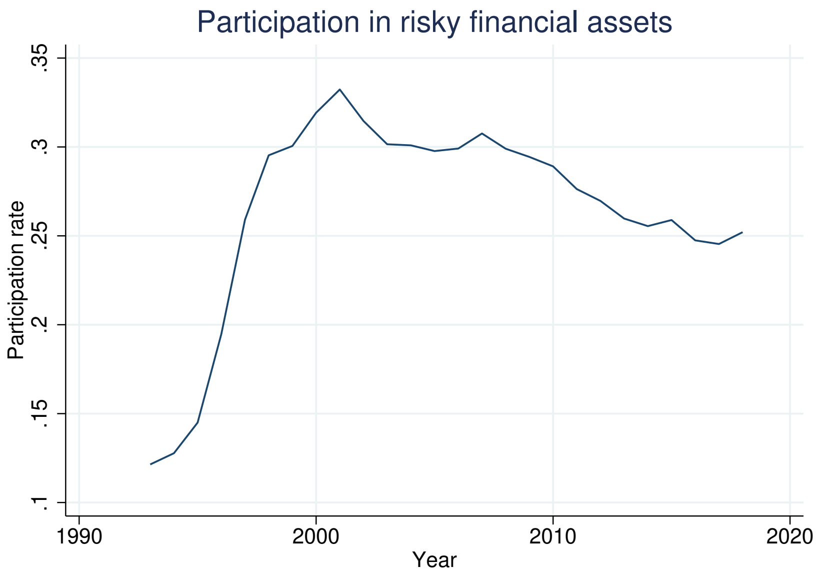

Note: this figure plots the participation rate in the stock market annually from 1993 to 2018.

FIGURE D.2: Stock market participation rates over time by asset class

(A) Mutual funds

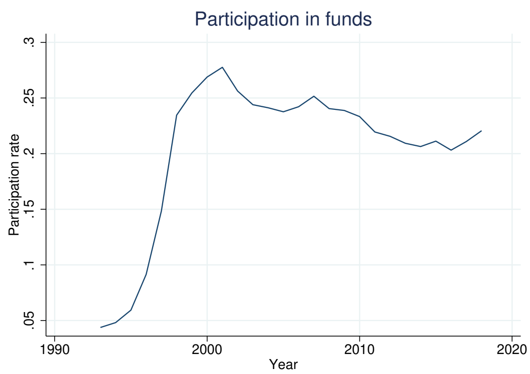

(B) Directly-held stocks

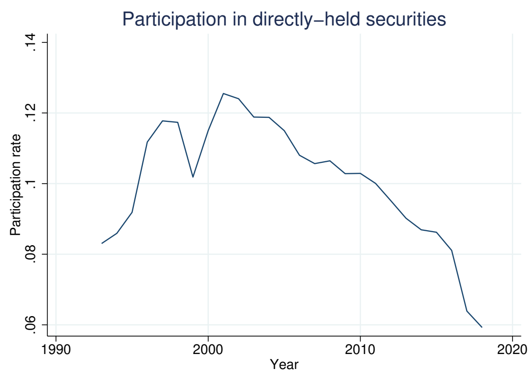

Note: this figure plots the participation rate in the stock market by asset class annually from 1993-2018. The left panel shows the participation rate in mutual funds, while the right panel is for directly-held stocks.

FIGURE D.3: Entry and exit rates over time

Note: this figure plots the entry and exit rates for stock market participation. The entry rate in year t is the proportion of non-participants in year $t-1$ who enter in year t. The exit rate in year t is the proportion of participants in year t who leave the stock market in year t.

FIGURE D.4: Proportion of participation spells ending within 1, 2 and 3 years

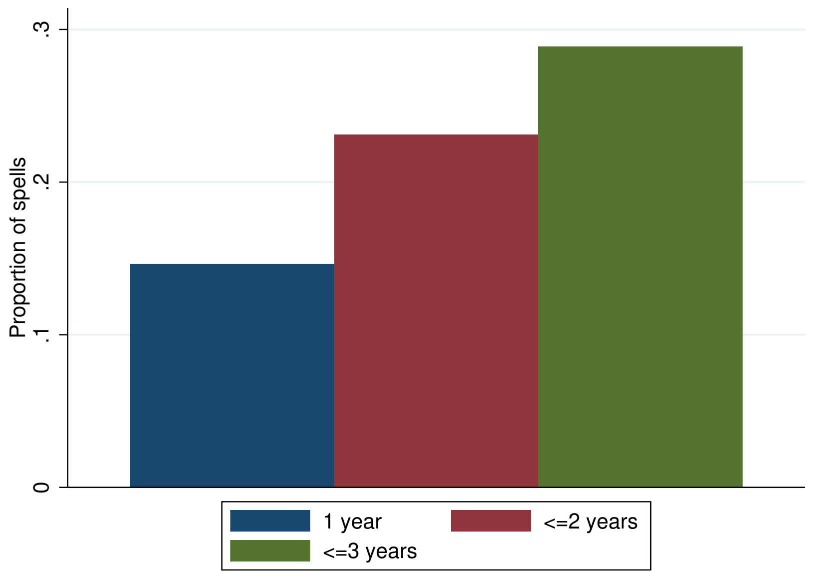

Note: this figure plots the proportion of all participation spells ending within 1, 2 and 3 years using participation spells beginning by 2017, 2016 and 2015 respectively.

FIGURE D.5: Spell length distribution at the household level

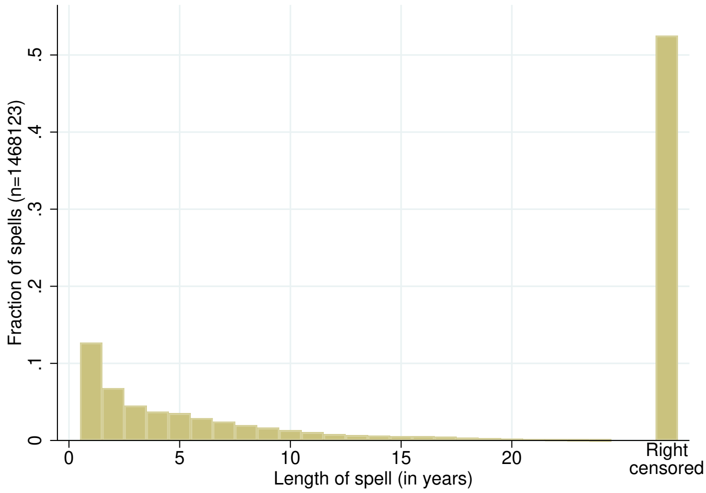

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data based on the household-level balance sheet. We take all spells beginning at any point from 1994-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells (n=1.5m) belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

FIGURE D.6: Spell length distribution (robustness to gifts/inheritance)

(A) No gift above 10,000 NOK

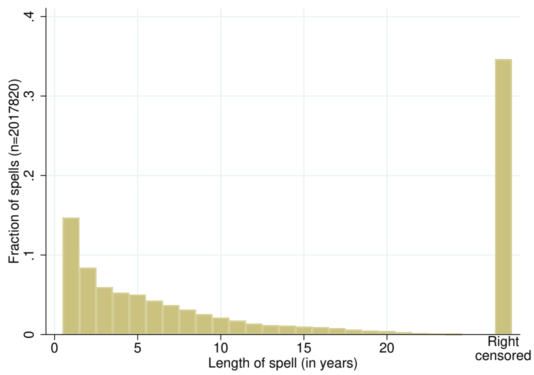

(B) No (grand)parent death

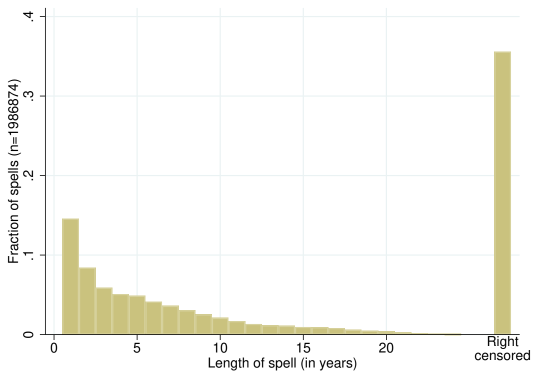

(C) No (grand)parent participation

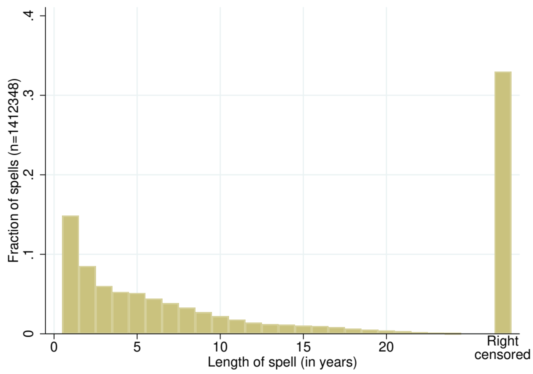

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data for different subsamples intended to deal with concerns that short spells are driven by gifts and inheritances. For all panels, we take spells beginning at any point from 1994-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored. Panel (A) excludes all individuals who receive a gift or inheritance above 10,000 NOK (based on tax records) in the year of or before entry. Panel (B) excludes all entrants who experience the death of a parent or grandparent in the year of or before entry. Panel (C) excludes all entrants for whom a parent or grandparents was participating in the year of or before entry.

FIGURE D.7: Spell length distribution excluding employee stocks

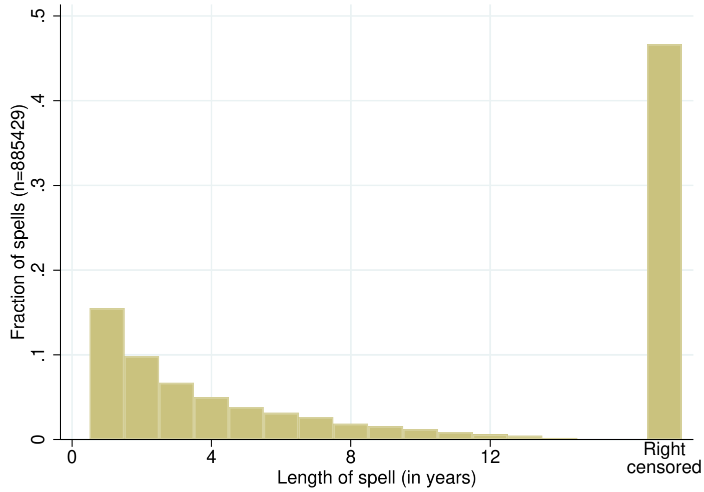

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data excluding entrants who hold stocks in the company they work for. We take all spells beginning at any point from 2004-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

FIGURE D.8: Spell length distribution excluding small investors

(A) Invest > \$100

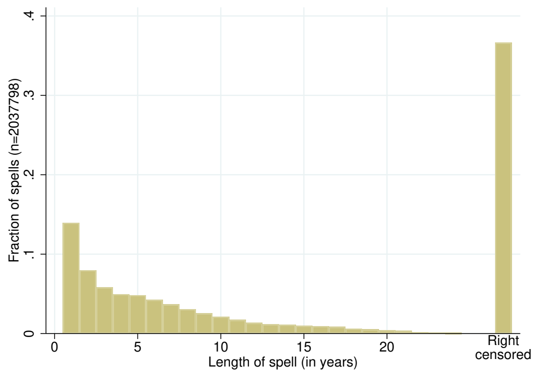

(B) Invest > \$1000

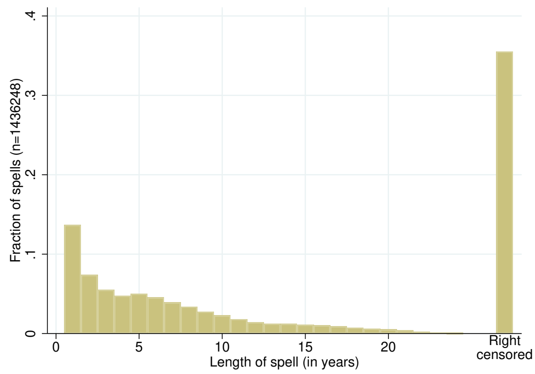

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data excluding entrants who invest a "small" amount of money at the point of entry. The left panel only uses individuals who invest at least \$100 at the point of entry, while the right panel requires an investment of at least \$1,000. hold stocks in the company they work for. For both panels, we take spells beginning at any point from 1994-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

FIGURE D.9: Spell length distribution by observable characteristics

(A) Income

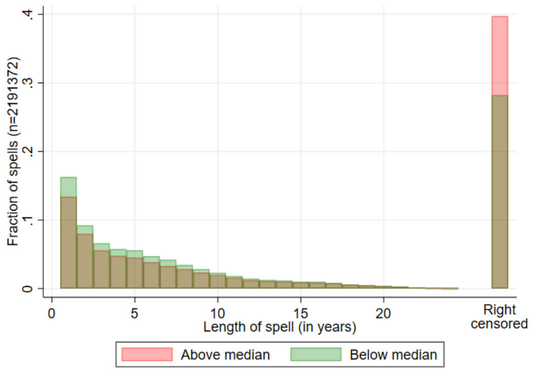

(B) Wealth

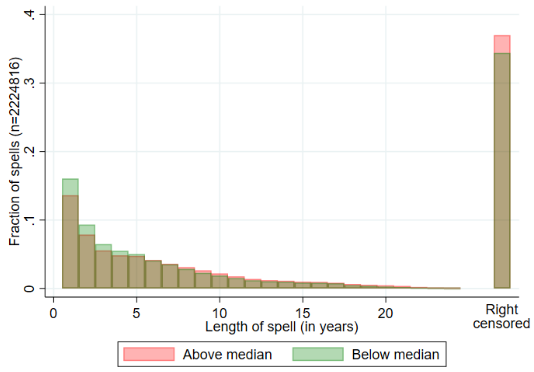

(C) Education

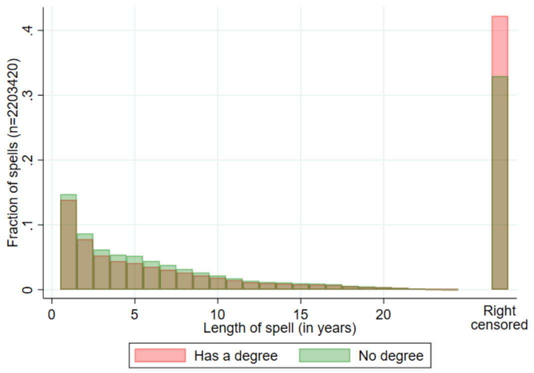

(D) Gender

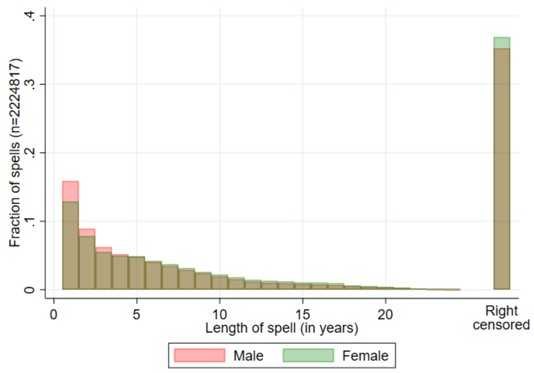

(E) By asset class

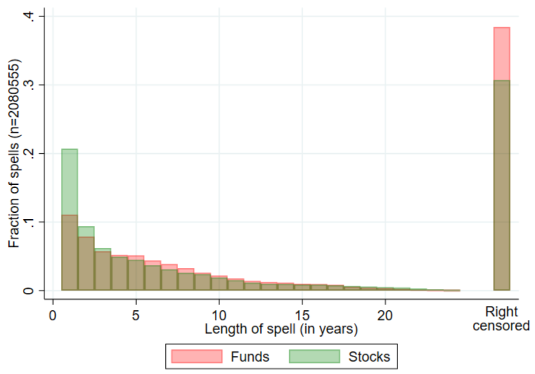

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data for different observable characteristics. Panels (A) and (B) show the distributions based on income and wealth respectively (below and above median). Panel (C) looks at the distributions for those with and without a college degree, while Panel (D) plots the histogram by gender. Panel (E) looks at individuals who enter into mutual funds vs. directly-held stocks. For this panel, we exclude those entrants who choose to invest in both at the point of entry. For all panels, we take all spells beginning at any point from 1994-2015. The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

FIGURE D.10: Distribution of re-entry times (robustness to gifts/inheritance)

(A) No gift above 10,000 NOK

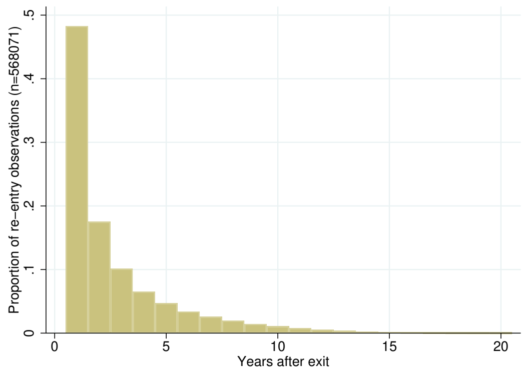

(B) No (grand)parent death

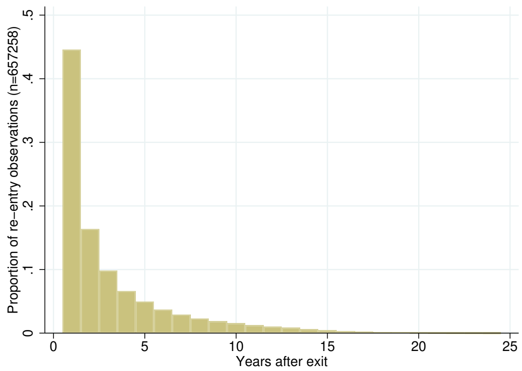

(C) No (grand)parent participation

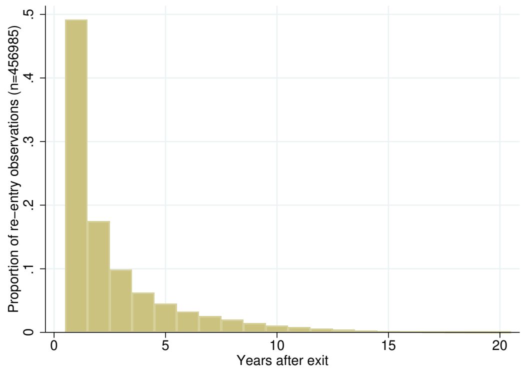

Note: this histogram plots the distribution of re-entry times in the Norwegian data for different subsamples intended to deal with concerns that short spells are driven by gifts and inheritances. The x-axis gives the re-entry time (in years) and the y-axis shows the proportion of re-entry observations belonging to a particular length. Panel (A) excludes all re-entrants who receive a gift or inheritance above 10,000 NOK (based on tax records) in the year of or before re-entry. Panel (B) excludes all re-entrants who experience the death of a parent or grandparent in the year of or before re-entry. Panel (C) excludes all re-entrants for whom a parent or grandparents was participating in the year of or before re-entry.

FIGURE D.11: Distribution of re-entry times (excluding employee stocks)

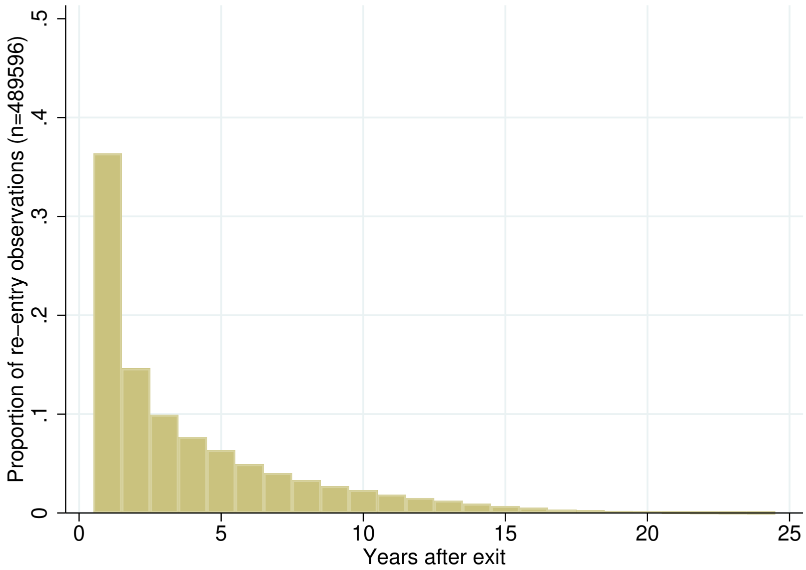

Note: this histogram plots the distribution of re-entry times in the Norwegian data excluding re-entrants who hold stocks in the company they work for. The x-axis gives the re-entry time (in years) and the y-axis shows the proportion of re-entry observations belonging to a particular length. As the Shareholder Registry data is only available from 2004, we only consider re-entry observations where the year of re-entry is no earlier than 2004.

FIGURE D.12: Prevalence of short spells over time

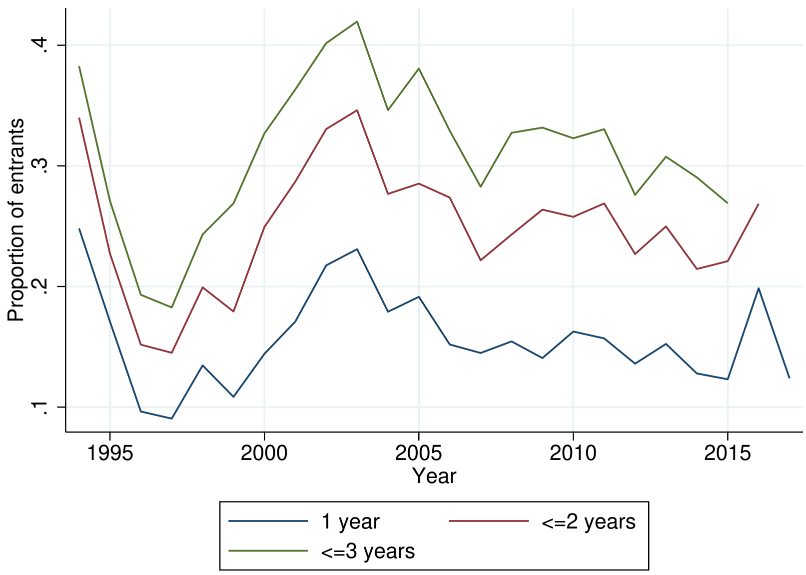

Note: this figure plots the proportion of entrants of a given year who exit within the next 1, 2 or 3 years.

FIGURE D.13: Performance of exiters by spell length

(A) Report gains

(B) Report losses

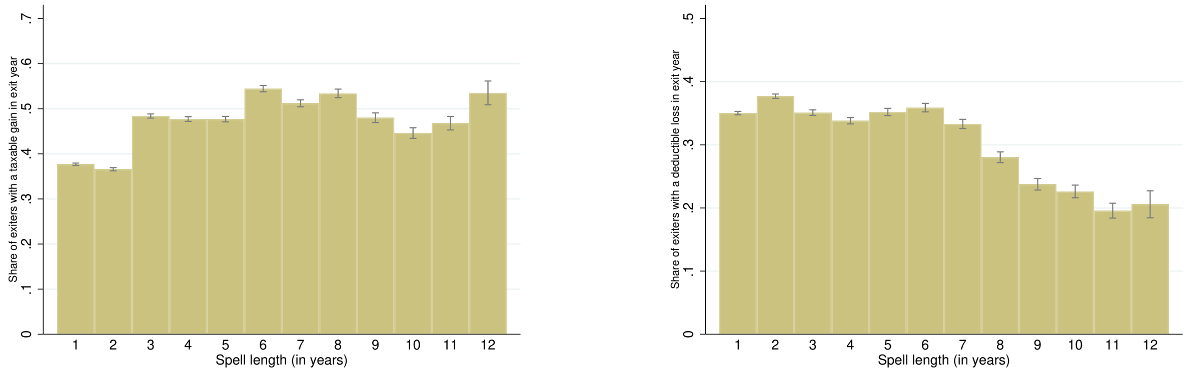

Note: this figure shows the performance of exiters by spell length based on records of taxable gains and tax-deductible losses in the income tax data. In panel (A), we plot the proportion of exiters of a given spell length reporting some gains (irrespective of losses) from the sale of stocks and funds (gains are computed as the sum of items TR 3.1.8, TR 3.1.9 and TR 3.1.10 in the tax records) in their exit year. In panel (B), we plot the proportion of exiters reporting some losses (irrespective of gains). Losses are computed as the sum of items TR 3.3.8, TR 3.3.9 and TR 3.3.10 in the tax records. We use exiters who enter from 2006 onwards in these plots.

FIGURE D.14: Prevalence of liquidity shocks in exit year by spell length

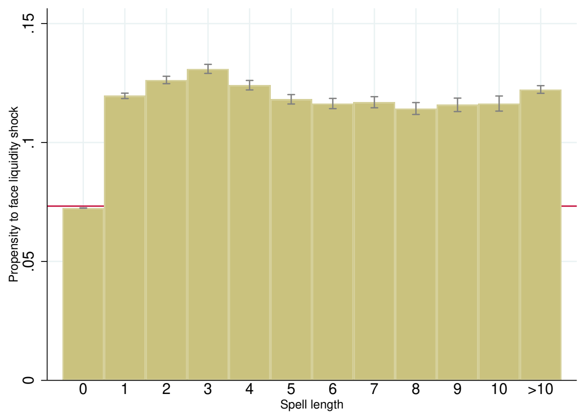

Note: this figure shows the proportion of exiters of different spell lengths experiencing at least one of three potential liquidity needs in the exit year. The three shocks considered are buying a house (observed in housing transactions data), divorce and unemployment (inferred through receipt of unemployment benefits). The far left bar (spell length of zero) gives the prevalence of liquidity shocks over non-exit observations (i.e. non-participants and continuing participants). The far right bar groups all exiters of spell lengths above 10 years. The red line gives the unconditional probability of experiencing a shock in the full population. 95% confidence intervals are shown.

FIGURE D.15: Participation in private pensions over time

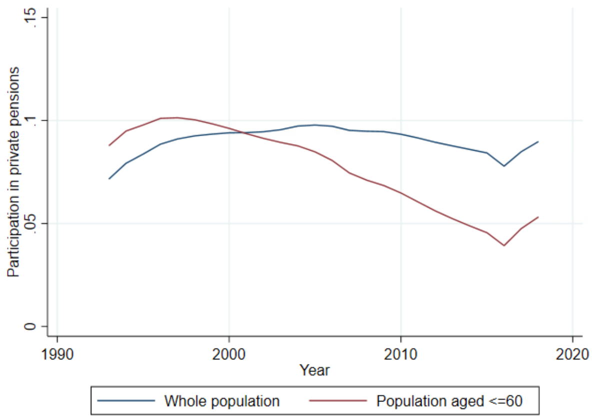

Note: this figure plots a time series of participation in private pensions over time. The blue line gives the participation rate for the whole population, while the red line restricts attention to those aged 60 or under. An individual is said to be participating in private pensions in a given year t if they put money into a private pension either in the current year or in a past year. Participation has occurred if either of the following two items in the tax records is non-zero: item 3.3.5, which records deductible payments to an Individual Pension Scheme (IPS), or item 4.5.1, which gives capital in an Individual Pension Account (IPA).

FIGURE D.16: Prevalence of private pensions amongst exiters by spell length

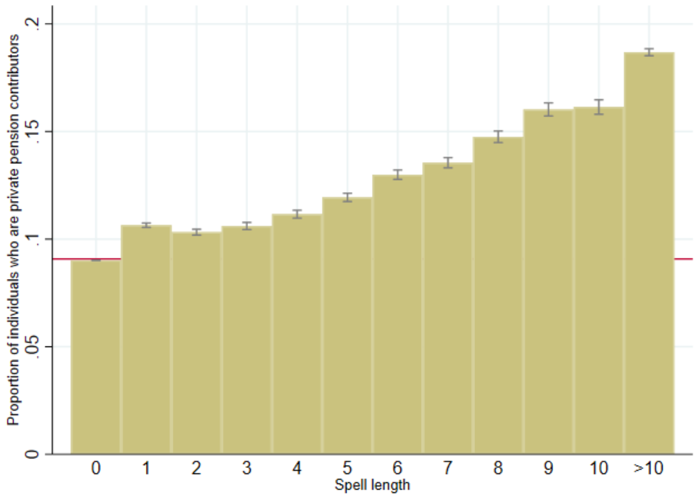

Note: this figure shows the proportion of exiters of different spell lengths participating in private pension accounts as of their exit year. An individual is said to be participating in private pensions in their exit year if they put money into a private pension either in the current year or in a past year. Participation has occurred if either of the following two items in the tax records is non-zero: item 3.3.5, which records deductible payments to an Individual Pension Scheme (IPS), or item 4.5.1, which gives capital in an Individual Pension Account (IPA). The far left bar (spell length of zero) gives the prevalence of private pensions shocks over non-exit observations (i.e. non-participants and continuing participants). The far right bar groups all exiters of spell lengths above 10 years. The red line gives the unconditional probability of experiencing a shock in the full population. 95% confidence intervals are shown.

FIGURE D.17: Spell length distribution excluding individuals with a private pension account

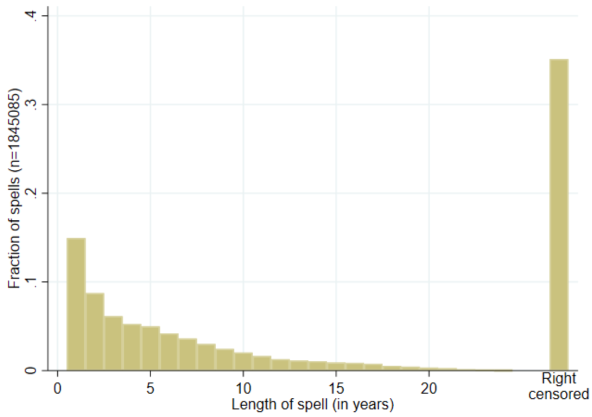

Note: this histogram plots the proportion of spells of different lengths in the Norwegian data excluding individuals who at any point in the sample hold a private pension account. Participation has occurred if either of the following two items in the tax records is non-zero: item 3.3.5, which records deductible payments to an Individual Pension Scheme (IPS), or item 4.5.1, which gives capital in an Individual Pension Account (IPA). The x-axis gives the spell length (in years) and the y-axis shows the proportion of spells belonging to a particular spell length. The far-right bar gives the proportion of these spells that are right-censored.

FIGURE D.18: Simulated mean consumption, wealth and labour income in the life-cycle model

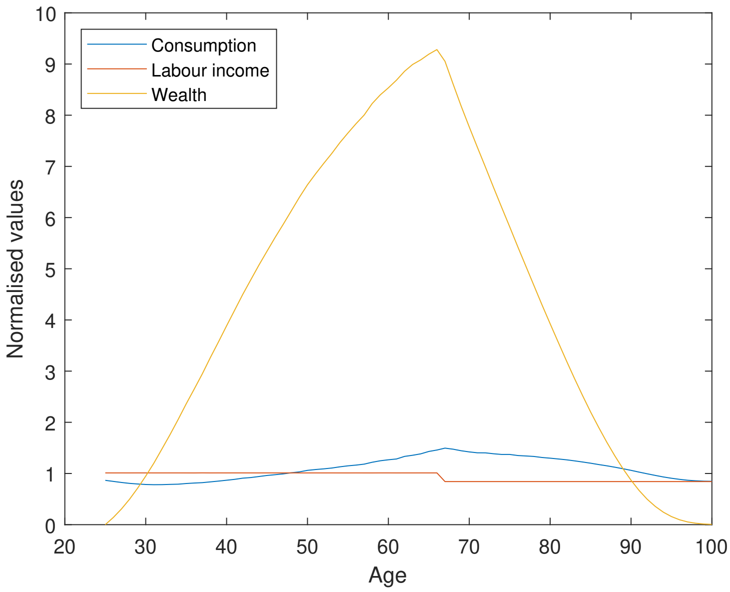

Note: this figure plots the simulated mean *normalised* consumption, wealth and labour income from the life-cycle model in Section 5.1. Variables are normalised with respect to the current level of permanent income. This is based on a simulation of 10,000 individuals.

FIGURE D.19: Simulated hazard function for exit

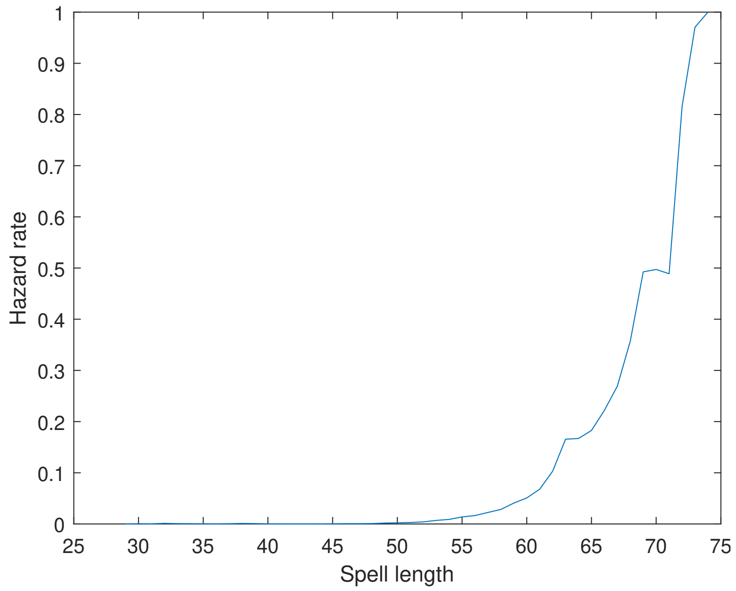

Note: this figure plots the hazard function for exit from participation. This is based on a simulation of 10,000 individuals.

FIGURE D.20: Simulated re-entry times

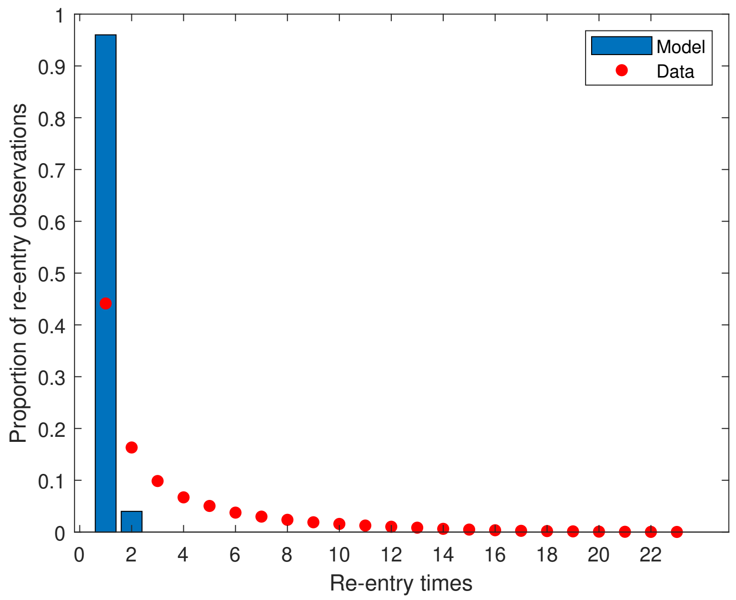

Note: this figure plots the simulated re-entry times. This is based on a simulation of 10,000 individuals.

FIGURE D.21: Spell length distribution by ability group

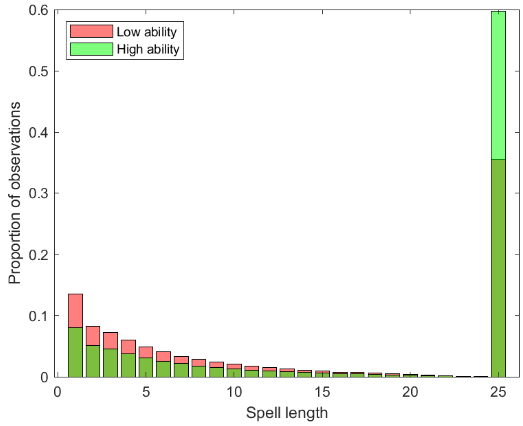

Note: this figure plots the distribution of spell lengths by ability group in the simulated model. The parameter values are given in Table 5. The model is simulated for $T = 25$ with N_t new entrants in each period, where N_t is given by the number of entrants in the Norwegian data. The red bars show the distribution for low ability individuals, while the green bars give the distribution for high ability individuals.