Heterogeneity and Persistence in Returns to Wealth *

Andreas Fagereng† Luigi Guiso‡ Davide Malacrino§ Luigi Pistaferri ¶

First Version: December 2015 This version: March 11, 2016

Preliminary and incomplete: please do not quote

Abstract: Heterogeneity, persistence, and correlation with wealth of individuals' returns to savings are key to explain wealth concentration at the top. We provide the first systematic analysis of the properties of individual returns to wealth using twenty years of population data for Norway. The data consist of information on income from capital and asset values from administrative tax records. We establish a number of novel results. First, in a given cross section individuals earn markedly different returns on their assets, with a range of over 500 basis points between the 10-th and the 90-th percentile. Second, heterogeneity in returns does not arise because of differences in how people allocate their wealth between safe and risky assets. Indeed, returns are heterogeneous even within asset classes. Third, returns are positively correlated with wealth. Fourth, individual returns have a persistent component that explains almost 20% of the variation. Fifth, permanent differences in individual returns reflect differences in initial wealth, level and type of education, as well as access to private equity investment. Sixth, the fixed component of the return to wealth is also (mildly) correlated across generations. Finally, there is assortative mating in returns. We discuss the implications of these findings for the debate about the drivers of wealth concentration, its measurement, and the relation between income and wealth inequality.

Keywords: Wealth inequality, returns to wealth, heterogeneity, intergenerational mobility, assortative mating.

JEL codes: E13, E21, E24

---

*We thank Alberto Bisin, Fatih Guvenen, Daniele Massacci, Andrea Pozzi and Anton Tsoy for very fruiful discussions. We are grateful to seminar participants at the 2016 ASSA meetings, the University of Chicago and the European University Institute.

†Statistics Norway, faa@ssb.no

‡Einaudi Institute for Economics and Finance (EIEF) and CEPR, guiso@tin.it

§Stanford University, malacrino@stanford.edu

¶Stanford University and NBER, pistaferri@stanford.edu

1 Introduction

Over time and across countries, the wealth distribution appears extremely skewed and with a long right tail: a small fraction of the population controls a large share of the economy's wealth. In the US, for example, the top 0.1% hold about 20% of the economy's net worth. Moreover, tail inequality has tripled in little more than three decades (Saez and Zucman, 2016). These striking regularities led Vilfredo Pareto to introduce a statistical distribution, now bearing his name, to model long-tailed economic phenomena and to theorize about possible socio-economic factors that might generate them.

After many years of debates, what produces the long tail of the wealth distribution and its extreme skewness is still a topic of intense research. Recent studies by Kopczuk (2015), Piketty and Saez (2003) and Saez and Zucman (2016) have revived the debate about the determinants of wealth concentration. One strand of literature, started by Aiyagari (1994) focused on the role played by idiosyncratic and uninsurable labor income (i.e., human capital) risk, leading households to accumulate assets for precautionary reasons. However, while idiosyncratic shocks to labor income can explain a non-negligible amount of wealth inequality, they are not enough to reproduce the extent of wealth inequality we measure in the data. In particular, they fail to account for the fat right tail of the distribution. The reason is simply that at higher levels of wealth the incentives to further accumulate precautionary assets fade away. Several authors have explored various additional mechanism that can potentially produce inequality and persistence in the distribution of wealth, such as non-homothetic preferences for bequests, heterogeneity in entrepreneurial talent, extreme skewness in the distribution of earnings for top earners, and heterogeneity in discount rates (e.g. Cagetti and De Nardi, 2006; Quadrini, 1999, 2000; Krueger and Kindermann, 2014; Krusell and Smith, 1998;). These factors generate more wealth inequality than in a simple Bewley-Aiyagari framework and come close in some instances to match the extent of inequality, concentration and persistence in wealth observed in the data. However, they are based on assumptions that are either counterfactual (such as the degree of skewness in the earnings distribution), or hard to evaluate empirically (such as the extent of discount rate heterogeneity).

More promising avenues has been recently explored by Benhabib et al. (2016) and Gabaix et al. (2016). Benhabib et al. (2016) build a theoretical model of wealth accumulation and intergenerational transmission of wealth that has some of the features of the literature cited above (such as a non-homothetic bequest motive and idiosyncratic returns to human capital), but departs from it by also allowing for heterogeneous stochastic returns to wealth. This latter feature and the properties of the distribution of returns to wealth turns out to be key for explaining wealth concentration at the top. Provided that the return to wealth has an important individual-specific component that persists over time and (to some extent) across generations, the model can generate a steady state distribution of wealth with a thick right tail populated by those who have been lucky (or expert) enough to get a repeated sequence of high returns to wealth. One attractive feature of this model is that it can naturally account for one important property of the data: the wealthy are typically the entrepreneurs, which in their model are individuals with access to an individual specific technology for generating idiosyncratic returns to wealth. Gabaix et al. (2016) show that while the Benhabib et al.'s model can explain the long thick tail of the wealth distribution, it cannot explain the speed of changes in tail inequality that we observe in the data. They suggest that one way to capture the latter is to allow for *type dependence* in the growth rate of wealth, i.e., high-wealth individuals have faster random growth rates of wealth than low-wealth individuals. Since the growth rate of wealth coincides with the return to wealth (absent savings or borrowing), the Gabaix et al.'s model requires that returns to wealth are positively correlated with the level of wealth.

But how much heterogeneity in returns to wealth is there in the data? Do returns to wealth persist over time within a generation as required by the Benhabib et al (2011) model? Do they persist across generations, and if so, by how much? Are returns and their heterogeneity correlated with wealth, as required by the model of Gabaix et al. (2016) designed to explain the fast increase in tail inequality? More generally, what are the empirical properties of the returns to wealth? In this paper we aim at answering these questions and thus provide background information useful to assess whether idiosyncratic returns to wealth can help explaining the empirical distribution of wealth and, in particular, whether it is able to generate its thick right tail. To this purpose, we rely on administrative tax records from Norway. These data contain information on both income from capital and wealth stocks. Wealth data include information on the value of all assets, real and financial, owned by each taxpayer in Norway.

As we will discuss, measurement error and underreporting of wealth information are unlikely to be a problem. This is because wealth data are generally collected through third parties (i.e., information provided by financial intermediaries) rather than being self-reported. Furthermore, the data have universal coverage, implying that there is exhaustive information on the assets owned by all individuals, including those at the very top of the wealth distribution. This is critical for a study of our sort, as leaving out the wealthy could potentially seriously understate the extent of heterogeneity in returns to wealth, particularly if returns are correlated with wealth and if the extent of heterogeneity also varies with wealth. Most importantly, the data have an extraordinary long panel dimension, covering 20 years – from 1994 to 2013 – and various business cycles. This allows us to study within-person persistence in returns. In addition, because over a 20-year period (some) generations overlap and because we can identify parents and children, one can also study intergenerational persistence in returns to wealth. Finally, since we observe individuals before they marry or form joint tax units, we can study whether returns to wealth persist across marital statuses and whether this reflects a form of assortative mating on returns to wealth.

We find that returns to wealth exhibit substantial heterogeneity. For example, in the last year for which we have data (2013) the (unweighted) average (median) return on overall wealth is 3.2% (2%), but it varies significantly across households. In particular, the standard deviation of the returns is 5%. When looking at returns from safe and risky assets separately, there are also large differences. In 2013 the average return to risky assets is 5.8%, more than double the return on safe assets, 2.5%. However, the standard deviation of the former (23%) is one order of magnitude larger than the standard deviation of the latter (3%). Hence, we find returns heterogeneity even when we focus on safe assets, although the dominant source of returns heterogeneity admittedly originates from heterogeneity in risky assets. Furthermore, heterogeneity in returns is not simply the reflection of differences in portfolio allocations between risky and safe assets mirroring heterogeneity in risk aversion. Even conditioning on the share of risky assets in portfolio, heterogeneity in returns is large and increases with the level of wealth. This result is confirmed even when looking at individuals with no private equity component in their risky asset portfolio. Another remarkable finding is that asset returns increase with wealth. In 2013, the difference between the median return for people in the 90th and 10th percentiles of the wealth distribution is 180 basis point.

In a given year, heterogeneity in returns to wealth may arise both from idiosyncratic transitory variations as well as for a persistent component in returns to wealth. To identify the latter we estimate a panel data statistical model for the returns to wealth that includes an individual fixed effect. To capture the heterogeneity that is explained by observable factors, we add controls such as portfolio composition, occupation, etc. The individual fixed effect captures the component of unobserved heterogeneity that persists over time. Finally, there is a component of heterogeneity that is unobserved but unsystematic (good/bad luck, etc.). We find that observable characteristics, alone, explain roughly 11% of the variability in returns to wealth; the inclusion of individual fixed effects increases explained variability substantially to 25%. The distribution of these fixed effects is itself quite disperse, with a standard deviation of 3 percentage points and a 90th-10th percentile difference of 6 percentage points. Interestingly, the distribution of the returns fixed effects is both more spread out and shifted to the right for people at the top of the wealth distribution compared to people at the bottom; for firm-owners compared to non-owners; and for people with higher levels of financial literacy

Having established the presence of significant systematic heterogeneity in asset returns over the life cycle, we turn to analyzing intergenerational and intramarital persistence in asset returns. We find that both the return from wealth and the fixed component of it are correlated intergenerationally, although there is rapid and strong mean reversion. Interestingly, the association between a child's asset return and the parent's asset return, while positive for a good range of the distribution, turns negative when the parent's return is above the 80th percentile. In other words, the children of individuals who were able to achieve very high returns from wealth have returns that, while still above average, revert more quickly to the mean.

We also find evidence of assortative mating in returns *conditional* on assortative mating in wealth. High-return singles tend to marry with individuals who also earn above-average returns. Interestingly, we find that post-marriage household returns mostly resemble the pre-marriage return of the highest-return spouse. However, the lowest-return spouse also plays a role, providing a rationale for assortative mating on returns to wealth. Interestingly, we find that the weight played by the highest-return spouse is higher if that spouse is the male.

As far as we know, this is the first paper that provides systematic evidence on individual returns to wealth over the entire wealth distribution, characterizes their properties, documenting the extent of cross sectional heterogeneity (both observed and unobserved), its correlation with the level of wealth, and its long-run persistence. Bach et al. (2016) perform an exercise similar to ours in spirit, but our paper differs from their in several respects. First, we have access to longer data than they do, allowing us to study returns persistence. Second, we observe all components of financial wealth, including private equity which is the dominant source of wealth for the very top fractiles of the wealth distribution. Third, we can study heterogeneity and persistence in returns to wealth over and above the intra-generational dimension that they focus on. Indeed, our paper is the first to provide systematic evidence on persistence in returns across generations and across marital statuses. This feature is critical for explaining the long thick tail in the wealth distribution. Benhabib, Bisin and Luo (2015) is the only paper we know of that estimates the extent of persistent heterogeneity and intergenerational persistence in returns to wealth. Their estimates are obtained from structural estimation of a life cycle model of wealth accumulation, calibrated on US data. Their paper finds evidence of both persistent heterogeneity and intergenerational persistence in returns to wealth and argues that both are critical for explaining wealth concentration at the top and mobility in wealth. However, the paper imposes lack of correlation between asset returns and the level of wealth, which instead we find is an important feature of the data. We also find that heterogeneity in returns varies over time. While heterogeneity in returns matters for explaining the level of wealth inequality at the top, variation over time in heterogeneity may matter for explaining variation in wealth inequality over time. With the exception of Gabaix et al. (2015), most papers have focused on explaining the distribution of wealth (or income) at a point in time assuming the economy is in steady state. This theoretical debate lags behind the empirical one that has shifted from measuring the extent of inequality at a point in time to documenting significant dynamics in inequality either in income (Saez and Piketty, 2003) or in wealth (e.g., Saez and Zucman 2016). None of the theoretical papers on wealth inequality has studied the implications of assortative mating in returns to wealth for wealth inequality and mobility, probably because while assortative mating on income is widely documented, our is the first study to show that people sort not only on wealth but also on returns to wealth.

The rest of the paper proceeds as follows. In Section 2 we review the literature. In Section 3 we present our data sources and discuss how we measure returns to wealth. Section 4 documents the extent of heterogeneity in returns and how returns on wealth correlate with wealth holdings. In Section 5 we discuss our empirical model of individual returns, showing how we identify persistent heterogeneity, presents the results and shows evidence of persistent heterogeneity within generations and persistence in returns across generations. Section 6 discusses some implications of heterogeneity in returns for the wealth inequality debate, while Section 7 concludes.

2 Heterogeneity in returns and the distribution of wealth

Absent sources of heterogeneity in saving propensities or sources of income other than labor, the distribution of wealth should inherit the properties of the distribution of earnings. Hence, if the distribution of labor incomes has a fat tail, the wealth distribution should mirror that feature. Yet, wealth seems to be uniformly more unequally distributed than income and realistic calibrations of heterogeneity in earnings that produce significant wealth inequality (as in Castaneda, Dìas-Giménez, and Rıos-Rull, 2003 and Krueger and Kindermann, 2014) do not seem to be able to account for the fatter tail in the distribution of wealth. For instance, while the calibrated model of Krueger and Kindermann (2014) gets close to matching the distribution of wealth in the US, it requires that the top 0.25% of income earners earn between 400 and 600 times more than the median earner. As Benhabib and Bisin (2015) notice, this is very far from what is observed in the data - where the ratio of the income of the top 0.1% percent to the median is only around 33. A similar argument applies to Castaneda et al. (2003).

One route, followed by Krusell and Smith (1998) has been to complement Bewley-Aiyagari models of earnings heterogeneity with heterogeneity in thriftiness, allowing individuals to differ in time discounting. Differences in thriftiness, together with heterogeneity in earnings, can considerably improve the match between the wealth distribution generated by the model and that in the data. Discount rate heterogeneity has a certain appeal because of its intuitive realism. On the other hand, discount rates are hard to observe and thus their heterogeneity difficult to assess. Hence, one has to impose and accept the heterogeneity that is needed to match the distribution of wealth without being able to validate it. Furthermore, discount rate heterogeneity seems to miss one important feature of the data: the high incidence of entrepreneurs at the top of the wealth distribution. Entrepreneurship is usually associated with higher risk tolerance and idiosyncratic risk (entrepreneurs tend to hold very high stakes in their own company - e.g. Heaton and Lucas, 2000; Moskowitz and Vissing-Jorgensen, 2002), rather than with higher than average discount rates. An alternative route followed in the attempt to match the thick tail in the distribution of wealth has been to allow explicitly for entrepreneurship and idiosyncratic returns to investment, as in Quadrini (2000) and Cagetti and De Nardi (2006). These papers show that a model that incorporates individual-specific technologies – i.e. entrepreneurs - can generate more wealth inequality than that produced by Bewley-Aiyagari models of earnings heterogeneity. In these models the driving factor that allows to match the observed wealth inequality is given by potentially high rates of return from entrepreneurial investment, coupled with borrowing constraints (which induce a selection of enterpreneurs among wealthy people to start with). Models of entrepreneurial idiosyncratic risk-taking have been developed more recently by Aoki and Nirei (2015) to explain the thick tail of the income distribution and its evolution over time, and by Benhabin, Bisin and Zhou (2016) using a more reduced form approach.

While idiosyncratic returns from entrepreneurship are one source of heterogeneity in returns to wealth that can help explain wealth concentration, heterogeneity in returns to wealth can arise from other sources. For example, Guvenen (2006) introduces return differentials by allowing all households to trade in a risk-free bond, but restricting one group of agents from accessing the stock market. One can view this model as capturing limited stock market participation and generating heterogeneity in returns to wealth between stockholders and non-stockholder. Guvenen (2009) shows that a calibrated version of this model can reproduce the differences in wealth holdings in the US between stockholders and non-stockholders.[1]

The heterogeneous stochastic returns approach to explain wealth concentration at the top has been more recently systematically developed and sharpened by Benhabib, Bisin and various coauthors in a sequence of contributions. Rather than focusing on the specific source of returns heterogeneity, they take the latter as given and study instead the consequences for the distribution of wealth of its presence and the properties that returns heterogeneity need to have for it to be able to account for the tail of the wealth distribution. In one key contribution, Benhabib, Bisin and Zhou (2011) consider an overlapping generation model where households differ both in returns to human capital and returns to wealth. Each household is endowed at birth with a rate of return on wealth and a return to human capital, drawn from independent distributions. Hence, there is persistence in returns to wealth (and human capital) within a generation. In addition, returns persist across generations and are independent of wealth. They show that in this model the stationary distribution of wealth has a closed form solution and is Pareto with a thick right tail. More importantly, it is the properties of the heterogeneity in returns that drive the thickness in the right tail of the wealth distribution, rather than the heterogeneity in returns to human capital. In other words, if return heterogeneity explains the tail of the wealth distribution, then the stochastic properties of labor income risk have no effect on the thickness of the tail of the wealth distribution (see their theorem 1). The latter is instead increasing in the degree of heterogeneity in asset returns. Benhabib and Bisin (2015) review the theoretical and empirical debate of the drivers of wealth inequality highlighting the specific role of returns heterogeneity. To assess quantitatively how far can heterogeneity in returns to wealth go in explaining the distribution of wealth and the degree of concentration in the tail (as well as the patterns of mobility in the wealth distribution), compared to other factors, they calibrate their overlapping generation model to US data. Besides heterogeneity in returns to wealth, the model allows also for heterogeneity in human capital and in

[1] Guvenen (2011) discusses the differential implications of his model of returns heterogeneity and models of discount heterogeneity as in Krusell and Smith (1998).

savings rates due to a bequest motive that varies with wealth. Benhabib, Bisin and Lou (2015) estimate the distribution of returns to wealth and its intergenerational persistence to match several moments of the US wealth distribution and the degree of intergenerational wealth mobility. They estimate average returns to wealth of 3.35% with a cross sectional standard deviation of 2.73%; intergenerational persistence in returns to wealth is positive and modest. Yet, even this amount of persistent heterogeneity is able to play a key role in matching the tails: indeed, the top 1% wealth share predicted by the model is almost identical to the equivalent moment in the data (33.6% in the data, 34.1% in the simulated model). Shutting down this channel alone by forcing returns to wealth to be the same across individuals would produce a top 1% share of only 5.7% and the overall wealth share of people above the 95th percentile to be 9.5% vis-à-vis an observed level of 60.3%. At the same time, wealth shares at the bottom of the distribution would be abnormally inflated. Returns heterogeneity appears thus a key factor for matching the empirical wealth distribution.

Gabaix et al. (2016) are interested not only in the amount of wealth concentration in the steady state, but also on the speed of the transition across steady states. They show that while the Benhabib et al.'s model can explain the long thick tail of the wealth distribution, it cannot explain the speed of changes in tail inequality that we observe in the data. They suggest that one way to capture the latter is to allow for *type dependence* in the growth rate of wealth, i.e., high-wealth individuals have faster random growth rates of wealth than low-wealth individuals. Since the growth rate of wealth coincides with the return to wealth (absent saving or borrowing), the Gabaix et al.'s model requires that returns to wealth are positively correlated with the level of wealth.

Despite the theoretical appeal, explanations for the level and the dynamics of wealth inequality and concentration based on a more sophisticated process for the returns to wealth suffer from some of the problems of the models that rely on heterogeneity in discount rates. How reasonable is the heterogeneity and persistence estimated in Benhabib, Bisin and Lou model (2015)? Is there a correlation between wealth and returns to wealth that is compatible with the speed of tail inequality observed in the data? Differently from individual discount rates, however, individual returns on wealth have the great advantage that they can be observed (though not easily). Yet, even if data can be retrieved, what needs to be documented is that returns on wealth have an individual component; that this component persists a lot across individuals of the same generation; that it correlates with wealth; and that it shows some intergenerational persistence. Documenting these facts requires much more than just observability. More generally, returns on wealth may show features that a calibrated exercise should account for to properly characterize the role of wealth heterogeneity as a driver of wealth inequality. The goal of this paper is to provide a systematic characterization of these properties.

3 Data sources and variable definitions

Our analysis employs several administrative registries provided by Statistics Norway, which we link through unique identifiers for individuals and households. In this section we discuss the broad features of the data; more technical details are in the Appendix. We start by using a rich longitudinal database that covers every Norwegian resident from 1967 to 2013. For each year, the database provides relevant socio-economic information (sex, age, marital status, educational attainment, income, and gross wealth) and geographical identifiers. For the period 1993-2013 - the one we focus on here - we can link this database with tax records containing individual information on asset holdings and liabilities (such as real estate, financial assets, private businesses, and debt), as well as a detailed account of the individual's sources of income (from labor and capital). The value of asset holdings and liabilities are measured as of December 31 of each year. While tax records data typically include information on income, they rarely (if ever) contain information on wealth. In Norway's case, this happens because of a wealth tax mandating taxpayers to report in their tax filing not only their incomes but also their asset holdings.

The data we assemble have several, noteworthy advantages over those available for most other countries, particularly for the purpose of our study. First, our income and wealth data cover all individuals in the population who are subject to the income and wealth tax, including people at the very top of the wealth distribution. Given the extreme concentration of wealth at the top, this is a key feature of the data.[2] In particular, steady-state wealth inequality is likely to be very sensitive to even small correlation between returns and wealth. Moreover, the degree of correlation may vary, as we will document, according to the level of wealth, and be higher at high levels of wealth. Hence, missing the top wealth may understate the degree of correlation. Furthermore, the extent of returns heterogeneity may differ across the wealth distribution. These features can only be captured if the data include people at the very top of the wealth distribution. Second, in our data set most components of income and wealth are reported by a third-party (e.g., employers, banks and financial intermediaries) and recorded without any top- or bottom-coding. Because of this,

[2]Wealth concentration in Norway is high. In 2012, the top 0.1% owned about 10% of all net worth in the economy. For comparison, in the US the top 0.1% owned about 22% of all the net worth in the economy.

these data do not suffer from the standard measurement error that characterizes household surveys, where individuals self-report income and assets components (as for instance in the US Survey of Consumer Finances) and confidentiality issues censor extreme asset holdings around certain thresholds. This attenuates concerns that the heterogeneity in measured returns to wealth may reflect systematic measurement error possibly correlated with individual attributes (such as wealth itself). Third, the Norwegian data have a very long longitudinal dimension: this is necessary to identify persistent heterogeneity in returns which is the focus of this paper and is a key determinant of wealth concentration. And because the data cover the whole relevant population, it is free from attrition, except the unavoidable one arising from mortality and out-of-country migration. Fourth, our data have information not only on listed stocks but also on private equity holdings. Because private equity holders have large stakes in private businesses, this feature is important for pinning down the extent of heterogeneity. And because, as we will document, stakes in private equity strongly increase with wealth, this feature is also important for understanding the correlation between wealth and returns. Finally, unique identifiers allow us to match spouses (and hence to construct measures of household wealth, as well as accounting for wealth changes induced by family formation and dissolution), and to match parents with their children. The latter feature, together with the long panel dimension of the data, is key to study intergenerational persistence in returns to wealth, which, in turn, may be an important determinant of wealth inequality in the tail (Benhabib, Bisin and Zhu, 2011; Benhabib, Bisin and Luo, 2015). Besides these unambiguos merits, our data have also some shortcomings: one, not surprising, is the measurement of the private equity another the calculation of capital gains. We discuss them below and suggest remedies.

In our main analysis we focus on returns to financial assets, which include bank deposits, bonds, stocks of listed companies, and shares in non-listed companies - i.e., private businesses - mutual funds, and money market funds.[3] Below we briefly describe the administrative tax records on wealth and income and how we construct

[3]The main component of wealth that is left out of our analysis is housing and the related returns. We leave housing out of the analysis for two reasons. First, a practical reason: housing wealth data before 2010 are incomplete. Second, a conceptual reason. Returns on owner-occupied housing, which are the main component of housing wealth for the bulk of the population, are given by the services they provide. Thus, the returns on owner occupied housing would have to be imputed. This would introduce measurement error most likely overstating wealth returns heterogeneity. Because housing returns are essentially uncorrelated with stock returns (Curcuru, Heaton and Lucas, 2009), our estimates provide a conservative measure of returns heterogeneity. On the other hand, leaving housing returns out of the picture is unlikely to bias the correlations between returns on wealth and the level of wealth. In fact, for the period 2010-13 (when housing data is complete and accurate), the correlation between financial wealth and total wealth (financial wealth + housing wealth - debt) ranges between 0.98 and 0.99.

our measure of wealth returns.

Details of the mapping between the capital income tax component and the specific asset category are in the Appendix.

3.1 Administrative wealth and capital income records

Norwegian households are subject to both income tax and a wealth tax. Each year they are required to report their incomes and complete information on wealth holdings to the tax authorities. Tax record data are available on an annual basis since 1993.[4] The collection of tax information is mostly through third-parties. In particular, employers, banks, brokers, insurance companies and any other financial intermediaries are obliged to send both to the individual and to the tax authorities, information on the value of the assets owned by the individual and administered by the employer or the financial intermediary, as well as information on the income earned on these assets. For traded assets the value reported is the market value. For an individual who holds no stocks, the tax authority pre-fills a tax form and sends it to the individual for approval; if the individual does not respond, the tax authority considers the information it has gathered as implicitly approved. In 2009, nearly 2 million individuals (60 percent of the Norwegian tax payers) belonged to this category. If the individuals or households own stocks, then they have to fill in the tax statement - including calculations of capital gains/losses and deduction claims. The statement is sent back to the tax authority which, as in the previous case, receives all the basic information from employers and intermediaries and can thus check its truthfulness and accuracy. Stockholders are treated differently because the government wants to save on the time necessary to fill in more complex tax statements. This procedure - particularly the fact that financial institutions supply information on their customer's financial assets directly to the tax authority – greatly reduces the scope for tax evasion, and thus non-reporting or under-reporting of assets holdings is likely to be negligible.

For the last ten years of sample the financial wealth data report information at the level of the single financial instrument. In particular, we can identify each single listed stock or bond in the investor portfolio. These data are analogous to those for Sweden available for the years from 1999 to 2007 and used by Calvet, Campbell, and Sodini (2007) and by Bach, Calvet and Sodini (2015).

[4] The individuals in a household are taxed jointly (i.e., married couples) for the purpose of wealth taxation, and separately for the income tax.

3.2 Wealth aggregates and returns to wealth

For our analysis we group assets into two broad categories, safe and risky assets (w^s and w^m, respectively), and map them with the corresponding values of capital income from the tax returns. We define the stock of safe assets as the sum of cash, bank deposits, treasuries, money market and bond mutual funds, bonds and outstanding claims, and receivables. The stock of risky assets is defined as the sum of the market value of listed stocks (held directly or indirectly through mutual funds, $w^{m,l}$) and the assessed value of shares in private businesses and other non-listed shares, $w^{m,u}$. For non-listed stocks, assessed values are the product of the equity share held in the firm and the value of the company as reported in the latter's tax returns. The assessed value excludes net present value calculation of the firm or goodwill. The tax authority also has control routines designed to identify firms that underreport their value. Medium- to large-size firms (with turnover above NOK 5m, or US $ 500k) are obliged to have their balance sheet reports audited by a professional auditing firm, reducing the scope for accounting misstatements. Total wealth is therefore:

$$w_{it} = w^s_{it} + w^{m,l}_{it} + w^{m,u}_{it}$$

Capital income y_{it} includes income earned on safe assets i_{it} (the sum of interest income on bank deposits and the like, other interest income, interest on loans to companies and the yield from insurance policies), dividends (from both public and private equity, d_{it}), and realized capital gains and losses from all equity (g_{it}). Because dividends and capital gains/losses on listed and private firms are, for tax purposes, reported jointly, we cannot compute separately the return from public equity and private equity. We hence observe:

$$y_{it} = i_{it} + d_{it} + g_{it}$$

Figure 1 shows the composition of the individual portfolio (i.e., shares of wealth in safe assets, listed stocks held either directly or indirectly through mutual funds, and the share in private businesses) for people in different parts of the wealth distribution. Safe assets clearly dominate the asset allocation of people below median wealth. Public equity (especially through mutual funds) gains weight among people above median and below the top 1%. The share in private business is strongly increasing with wealth above the 95-th percentile and carries a very large weight, close to 90%, for the top 0.01%. Because returns to private equity are largely idiosyncratic, given the strong correlation between exposure to private business and wealth, lack of information on private equity holdings is likely to grossly understate returns heterogeneity.

3.3 The measurement of returns to wealth

Consider an individual who invests a given amount of wealth w_{it} in a financial instrument that pays an annual return r_t. Suppose that the individual's portfolio is passive throughout the period, so that the investment delivers a flow of income $y_{it} = r_t w_{it}$. The individual return on wealth could thus be estimated as:

$$r_{it} = \frac{y_{it}}{w_{it}} \tag{1}$$

Our measure of return to wealth has to account for three data limitations. First, we only observe snapshots of people's assets at the end of each period, while observing the flow of income from capital throughout the period. Second, the value of private equity does not necessarily correspond to the underlying market value. Finally, we only observe capital gains or losses when they are realized (i.e., when assets are sold), not when they accrue economically.

We account for these three limitations using different adjustment procedures. Consider the first problem. If assets are accumulated or decumulated during the year, the income from capital will only reflect the part earned over the holding period before (after) the assets sales (purchases). The issue is most obvious in the case in which beginning-of-period wealth $w_{it} = 0$ but $y_{it} > 0$. To account for this problem, we define returns as the ratio of income from capital and the average stock of wealth at the beginning *and* end of year, i.e.:

$$r_{it}^{(1)} = \frac{y_{it}}{(w_{it} + w_{it+1})/2} \tag{2}$$

To see the importance of this adjustment, consider an individual who has beginning-of-period wealth $w_{it} = 0$ and after six months invests \$100 in a money market account at 10% interest rate, earning $y_{it} = 5$. End-of-period wealth is thus $w_{it+1} = 105$. The naive measure of return (1) is undefined. The adjusted return measure is instead $r_{it}^{(1)}$=9.52%, much closer to the actual 10% return. The adjusted measure works well also when people withdraw for consumption purposes. Consider an individual who has beginning-of-period wealth $w_{it} = 100$ invested in a 10% money market account. After 9 months, the individual withdraws and spends \$50, so that capital income is $y_{it} = 8.75$. End-of-period wealth is $w_{it+1} = 58.75$. The naive measure of return (1) would be lower than the actual one, 8.75%. The adjusted return measure is, instead, $r_{it}^{(1)}$=11%, much closer to the actual 10% return. We use this adjustment both when we compute the returns on safe assets, $r_{it}^{s(1)} = \frac{i_{it}}{(w_{it}^s + w_{it+1}^s)/2}$, as well as when we measure returns on risky assets, $r_{it}^{m(1)} = \frac{d_{it}+g_{it}}{(w_{it}^m + w_{it+1}^m)/2}$.

Our sample selection is also designed to reduce errors in the computation of returns. First, we focus on individuals aged 20 to 70. Second, we drop people with less than NOK 3000 in financial wealth (about US $500). These are typically transaction accounts with highly volatile beginning- and end-of-period reported stocks that tend to introduce large errors in computed returns.[5] Finally, we trim the distribution of returns in each year at the top and bottom 0.5%. These are conservative corrections that, if anything, reduce the extent of return heterogeneity.

Consider the second problem. Our measure of wealth from risky assets is the sum of market-valued wealth $w_{it}^{m,l}$ and assessed-value private equity $w_{it}^{m,u}$:

$$w_{it}^m = w_{it}^{m,l} + w_{it}^{m,u}$$

Neglicting capital gains/losses, our measure of returns to wealth is overstated if private equity firms assess the value of the company at lower prices than what they would get if they were to sell the firm. We account for this problem by obtaining an alternative measure of the value of private equity. We regress the market value of listed firms onto their book value and other observable characteristics, also available for private equity firms. We then impute the value of private equity using the regression coefficients, so that our measure of wealth becomes:

$$\hat{w}_{it}^m = w_{it}^{m,l} + \hat{w}_{it}^{m,u}$$

Hence, our second measure of return is

$$r_{it}^{(2)} = \frac{y_{it}}{(w_{it}^s + \hat{w}_{it}^m + w_{it+1}^s + \hat{w}_{it+1}^m)/2} \tag{3}$$

The third problem with our data is that we observe capital gains/losses when they are realized instead of when they accrue year by year. As we show in the Appendix

[5]For example, an individual with a (close to) zero balance (say $150) at the beginning of the year and a (close to) zero balance at the end of year (say $150), perhaps because of above average Christmas expenditures and average balance during the year of $ 3500 (30000 Kr), would earn a return of $70 if the interest rate is 2%. But the computed returns would be 70/150=50%. This is less likely to happen for large accounts.

this is not a serious issues if we are interested in measuring the average returns to wealth over the life cycle of an individual and we observe enough realizations of the capital gains. In fact the average return over a holding period of T years of an asset that is sold at T is the same whether the average return is computed using the annual return $R(t) = \frac{y_t}{P_t} + \frac{P_{t+1}}{P_t}$ with capital gains computed on an accrual basis, or when the annual return is $R(t) = \frac{y_t}{P_t}$ if $t < T$ and $R(t) = R(T) = \frac{y_T}{P_T} + \frac{P_{T+1}}{P_1}$ if $t = T$ as in our data.[6] On the other hand, if the measured return on a risky asset for individual i is $R_i(t) = \frac{y_t}{P_t} + I_i \frac{P_{t+1}}{P_1}$ where $I_i = 1$ if i sells the asset at t (and we observe the capital gain) and zero otherwise, then clearly at each point in time this induces some cross sectional heterogenity in measured returns because in each year only a fraction of the individuals realizes the capital gain. In the Appendix we show that this type of heterogeneity is nevertheles contained and can only explain a small fraction of the heterogenity in returns on risky assets that we measure in each year.

We also follow a more direct route to deal with unrealized capital gains. For unlisted stocks, we assume that capital gains reflect the increase in value of the stocks, i.e., $\Delta \hat{w}^{m,u}_{it+1}$. For public equity, we assigne the stock market's aggregate capital gains to investors on the basis of their beginning-of-period total stock market wealth. Define $M_t = \sum_{j=1}^{J} P_{jt} q_j$ the aggregate stock market value where P_{jt} is te price of stock j and q_j its quantity; let the aggregate capital gain be $G_t = \sum_{j=1}^{J} \Delta p_{jt+1} q_j$. The individual unrealized capital gain/loss from stockholding can hence be estimated as:

$$UCG_{it} = \frac{w^{m,l}_{it}}{M_t} G_t$$

And our final return measure is thus:

$$r^{(3)}_{it} = \frac{\Delta \hat{w}^{m,u}_{it+1} + UCG_{it} + y_{it} - g_{it}}{(w^s_{it} + \hat{w}^m_{it} + w^s_{it+1} + \hat{w}^m_{it+1})/2} \tag{4}$$

From now on, we focus mostly on the return to total wealth (2), which has the advantage of being based on information directly available from the tax records. In Appendix Section A.2 and A.3 we show the sensitivity of our main findings to adopting the alternative measures of returns (3) and (4).[7]

[6] In other words, provided we observe the realization of the capital gain at some point over the life cycle, the individual fixed effects in the return equation that we will estimate in Section 5 will correctly capture the average individual returns and thus the persistent heterogenity across individuals.

[7] All returns statistics we report are at the individual, not household level. This way we account for the fact that while households form and dissolve, individuals can be observed as they cycle through different marital arrangements. When individuals are single, the formulae above apply

3.4 Descriptive statistics

Table 1 shows summary statistics for our data. For simplicity, we report statistics for the last year in our dataset (2013) and, for comparison, summary statistics for 1994 in the Appendix. Overall, our 2013 sample includes almost 3 million households. In Panel A we report some basic demographic characteristics. The sample is well balanced between male and female household heads, and marital status (49% are married). Slightly more than 80% of individuals in the sample have at least a high school degree. Finally, 12% of individuals have a degree (college or high school) with a concentration in economics or business, which may be indicative of possessing above-average financial literacy. In Panel B we start digging into statistics describing wealth levels and composition. In 2013, almost half of the households in our sample had some risky assets in their portfolio. One in ten owned shares of a private business. Conditioning on having some assets invested in risky instruments, households invested on average 30% of their portfolio in those risky instruments. There is more concentration among private equity holders. Conditioning on owning a business, almost half of the wealth owned is in the business itself. The last five rows of Panel B provide information on wealth levels. Total financial assets are on average abut $85,000. As expected, the distribution is extremely skewed, with a median of about $20,000.

In the last panel of Table 1 we report summary statistics for the returns. In 2013, the average return on overall wealth was 3% (the median 2%), and the standard deviation 5%. The average return on risky assets (5.8%) exceeded substantially that on safe assets (2.5%). Statistics for the whole period 1994-2013 are qualitatively similar, although quantitatively the differences are enhanced by weighting the returns by portfolio values. For example, the average returns are 3.2% and 3.7%, respectively in the unweighted and value-weighted case. Similarly, the average returns from risky assets are 3.5% and 4.9% in the two cases. The larger difference in the value-weighted case is explained, as we shall see, by the positive correlation between returns and wealth levels.

without modifications. When individuals are married, we assume that spouses share household wealth and capital income equally. This is consistent with the Norwegian law that upon divorce family assest are split equally between the spouses. We first compute the return on household wealth, and then assign to each spouse this return and the per-capita household wealth.

4 Stylized facts about returns to wealth

In this section we establish a number of stylized facts about individual returns to wealth. In the next section we provide a formal framework to model returns to wealth that will help shed light on these stylized facts.

4.1 Returns to wealth are heterogenous

Figure 2 (top left panel) shows the cross sectional distribution of returns on total wealth in 2013, the last year of our sample. It makes clear that individuals earn markedly different returns. The average return on wealth is 3% with a standard deviation of 5% (Table 1, panel C)[8]. The median return is 2%, 100 basis points lower than the mean, implying a cross sectional distribution of returns on wealth that is significantly right-skewed. The difference between the median return at the 90th and the 10th percentiles is about 200 basis points. The other two panels show the distribution of returns for risky assets (top right panel) and safe assets (bottom panel).[9] The distribution of returns from risky assets is, as expected, more spread-out and with a higher mean.

To benchmark the extent of heterogeneity in returns one might expect from traditional household finance model, let's consider a standard Merton-Samuelson framework in which all investors have access to the same investment opportunities. In this model, investors choose the share of risky assets π_i as a function of market expceted excess returns, variance, and risk tolerance:

$$\pi_i = \frac{E(r_t^m - r_t^s)}{\gamma_i \sigma^2}$$

Heterogeneity is induced by differences in risk aversion, measured by γ_i. It follows that the individual realized return to total wealth is a weighted average of the risk-free rate and the market return:

$$r_{it} = r_t^s + \pi_i (r_t^m - r_t^s) \tag{5}$$

Equation (5) suggests that conditioning on having the *same* share of risky assets in portfolio, total returns on wealth should be similar across investors. That is, the

[8]The coefficient of variation in the Norwegian case is thus larger than that calibrated with US data by Benhabib et al. (2015), who find an average return of 3.4% with a cross sectional standard deviation of 2.7%. However, the calibration by Benhabib et al. (2015) refers to average individual returns over the lifecycle. We will discuss measures that are comparable to the ones they report in Section 5.

[9]Zero returns on risky assets correspond to individuals reporting no dividends and no realized capital gains or losses.

cross sectional standard deviation of returns, given π_i, should be close to zero. In Figure 3 we use again data for 2013. We allocate individuals to different bins defined by the share of their wealth held in risky assets (from 0 to 1), and within each bin we compute the cross-sectional standard deviation of the individual returns. Not only the standard deviation is not zero, but it also increases dramatically with the share of risky assets held in the portfolio. Interestingly, even at $\pi_i = 0$ (individuals own only safe assets), the standard deviation of returns is positive. Thus, while the composition of wealth (between risky and safe assets) does affect the extent of heterogeneity in the overall return to wealth, it is by no means the only driver (as we shall see more clearly in formal controlled regression, discussed in Section 5). Note that some of the heterogeneity in Figure 3 may come from holdings of private equity with very idiosyncratic, undiversified returns from entrepreneurship. We hence repeat the exercise focusing only on investors who do not own any private equity, i.e., individuals who only invest in safe assets and listed companies. The evidence is similar, although as expected the extent of heterogeneity is lower. Also as expected, this shows that there is much more risk involved in the holding of private equity wealth (see e.g,. Carrol, Moskowitz and Vissing-Jorgensen, 2002; Kartashova, 2014 and others).

Heterogeneity in returns is present in all years and its extent varies overtime. Figure 4 plots the cross sectional mean, median and standard deviation of returns on wealth for all sample years. Heterogeneity varies markedly over time with a cross sectional standard deviation of returns ranging between 0.08 in 2005 and just above 0.04 in 2009. Cross sectional heterogeneity of returns on total wealth does not depend on average returns. Figure 5 shows the patterns for returns on safe and risky assets. Heterogeneity on the latter covaries closely with average returns; heterogeneity in returns on risky assets is much higher, much more volatile and uncorrelated with average returns.

4.2 Returns covary with the level of wealth

Returns are correlated with the level of wealth. Figure 6, plots the median return on wealth for households in different percentiles of the wealth distribution using data for 2013. The differences in returns across wealth levels are large. Median returns for households at the 10th and 90th percentile of the wealth distribution are 0.7% and 2.6%, respectively. Hence, moving from the 10th to the 90th percentile of the wealth distribution the median return almost quadruples, suggesting that the correlation between returns and wealth holdings can potentially have large effects on wealth inequality.[10] How large requires new investigation. Indeed, recent calibrated models of wealth inequality by Benhabib, Bisin and Lou (2015) and Hubmer, Krusell and Smith (2015) allow for heterogeneity in returns to wealth but assume absence of correlation between returns and wealth. Such correlation is invoked by Gabaix et al. (2016) to explain the fast increase in tail inequality observed in many countries.

Correlation between returns and wealth may arise because of fixed entry costs in risky assets that preclude participation to low wealth households. This is indeed consistent with a large literature on limited participation costs (surveyed in Guiso and Sodini, 2012) and emphasized by Guvenen (2006) in the context of the wealth inequality debate. Moreover, it may simply reflect the fact that wealthy investors are more risk tolerant, have a riskier portfolio, and hence receive a premium for greater risk-taking. Finally, there are important economies of scale in wealth management. Recent work by Kapcerczyk et al. (2014) (building on earlier ideas by Arrow, 1987) suggests that wealthy investors are more "sophisticated" than retail investors, for example because they can access better information about where the market is heading, and hence reape higher returns on average.

The second panel shows that the positive correlation between returns and wealth holds both for risky as well as for safe assets. This rules out that the returns wealth correlation arises only because of participation costs in risky assets markets. Differences in returns on safe assets depending on wealth levels is instead consistent with differences in remuneration on deposits depending on amounts deposited; for instance, in 2008 this ranges between 4% per year for deposits less than 7,000 dollars to 6% for deposits larger than $37,000 (see Appendix, Figure A1).

The extent of heterogeneity also covaries with wealth as shown in Figure 7 which plots the cross sectional standard deviation of returns for each percentile of the wealth distribution using again 2013 as a reference year. In this year heterogeneity is relatively high at low levels of wealth and is fairly flat between the 20th and the 70th percentile, when it starts increasing more sharply, resulting in a U-shaped relation between the cross sectional standard deviation of returns and wealth. While the high-heterogeneity in returns at the bottom is not a feature of all years, the correlation at the top is (see Figure 9 below).

One way of summarizing the evidence above is to compute a measure of the Sharpe ratio at the individual level, using the 20 years in which the individual is potentially observed in our data. The individual Sharpe ratio is defined as:

[10] As noticed by Piketty (2014), "It is perfectly possible that wealthier people obtain higher average returns than less wealthy people.... It is easy to see that such a mechanism can automatically lead to a radical divergence in the distribution of capital".

$$S_i = \frac{\frac{\sum_{t=1}^{T} r_{it}}{T}}{\sqrt{\frac{\sum_{t=1}^{T} r_{it}^2}{T} - (\frac{\sum_{t=1}^{T} r_{it}}{T})^2}} \tag{6}$$

In Figure 8 we plot the average Sharpe ratio for each percentile of the wealth distribution in 1995 (the initial condition). Clearly, wealthier individuals reap higher returns for given amount of risk.

The extent of the correlation between returns and wealth is not specific to a given year. It appears as a defining feature of the data, although it does vary over time. To summarize these features in a simple way Figure 9 plots the median returns for households at selected percentiles of the wealth distribution over the 20 year period for which we have data. It shows very clearly that households in higher percentiles of the wealth distribution enjoy higher returns independently of year; it also shows that the difference in returns between high and low wealth levels varies considerably over the sample.

Fagereng et al., (2016) show that even a small positive correlation between returns and wealth can significantly overstate inequality measures when wealth is estimated by capitalizing income from tax returns as in Saez and Zucman (2016). We will discuss this evidence in Section 7.

4.3 Returns on wealth persist across generations

Our Norwegian data contain not only the individual identifier but also the family identifier. Hence it is possible to link individuals to their parents and/or their children. To focus on a sharper case, we look at fathers and sons. Because our sample covers several years, many individuals in our sample overlap with their parents. In principle, one would like to relate parents' variables and children's variables when they are of the same age. Unfortunately, our panel is not long enough to make this requirement practical. To control for the fact that parents and children are observed when they are at different points of their life cycles, we compute rank percentiles with respect to the birth cohort the individuals (father and son) belong to.

Figure 10 plots the average percentile of the son's return against the return percentile of his father. The figure shows a positive correlation, albeit the slope is far below the 45 degree line corresponding to perfect correlation. As the figure shows, the correlation tends to be stronger at intermediate levels of the parent's percentile; correlation turns negative - marking fast regression to the mean - for parents with very high returns on wealth. Some of the intergenerational correlation in returns may come from parents and children sharing a private business (or family firm). It is also possible that kids imitate the investment strategies of their parents, or that they inherit from their parents traits that matter for returns (such as preferences for risk or investment talent). Figure 10 shows correlations in realized returns; as argued by Benhabib, Bisin and Zhou (2011) models that aim at explaining the tail of the wealth distribution require wealth persistence in the generation-specific component. Furthermore, because wealth and returns are correlated, the correlation in returns across generations may well reflect correlation across generations in wealth. We defer a deeper inquiry of these issues to Section 5.

4.4 Returns heterogeneity entails assortative mating

Figure 11 documents our last stylized fact about returns to wealth: assortative mating in returns. To construct this picture, we focus on a sample of individuals who make a transition from singlehood to marriage at some point during our sample period. We start by computing the average return during the singlehood stage. In Norway people share a tax ID if they are married, or if they are cohabiting but sharing the care of a child. To avoid contamination induced by the fact that what appears as singlehood may be childless cohabitation, we experiment dropping two or four years before firstly observing individuals sharing a tax ID.

Figure 10 shows that pre-marriage average returns to wealth are positively correlated. This remains true regardless of whether we drop the two or four years preceding marriage (or the birth of a child to cohabiting couples). Clearly, the degree of assortative mating falls well below the 45 degree line of perfect sorting. As in the case of intergenerational correlation, assortative mating in returns may reflect assortative mating in attributes such as education or wages (which is well documented in the literature) or assortative mating on wealth (on which there is instead no evidence in the literature, mostly due to lack of data). Indeed, Figure 12 shows that spouses assort on their pre-marriage wealth. In Section 7 we address these issues and document that assortative mating in returns is a features that holds independently of assortative mating in wealth or other traits commonly identified as inducing sorting patterns in marriage.

5 Modeling returns to wealth

In this section we provide a formal statistical structure to model individual returns, characterize their heterogeneity and assess whether the heterogeneity that we have documented is just the reflection of idiosyncratic realizations that are quickly reversed or whether individuals that tend to earn higher returns tend to do so persistently over time. In other words, we investigate whether individual returns to wealth have a fixed effect component. Persistence in returns within and across generations, as argued by Benhabib, Bisin and Zhu (2011, 2015), is essential for heterogeneity to be able to explain the fat tail of the wealth distribution.

We specify a linear panel data regression model for wealth returns:

$$r_{igt} = X_{igt}^{'}\beta + u_{igt} \tag{7}$$

where r_{igt} denotes the return to wealth for individual i belonging to generation g in year t. X_{igt} is a vector of controls meant to capture predictable variation in returns due to individual observables, such as age, common shocks (time effects), etc.. To control for the risk induced by asset allocation, the vector X_{igt} includes the share of wealth held in risky assets. In a world where individuals are fully diversified, and thus invest in the same portfolio of risky securities, and have access to the same returns on safe assets, the portfolio return would be $r_{igt} = r_t^s + \pi_{igt}(r_t^m - r_t^s)$. Hence a regression of returns on time dummies and the individual risky assets share π_{it}, would absorb all the existing variation. If some individuals have access to private equity, as in Aoki and Nirei (2015) and Quadrini (2000), while continuing to invest in a fully diversified portfolio of listed stocks, return on wealth can be written as $r_{igt} = r_t^s + \pi_{igt}^m(r_t^m - r_t^s) + \pi_{igt}^{PE}(r_{igt}^{PE} - r_t^s)$, where now π_{igt}^m and π_{igt}^{PE} denote the share in listed stocks and private equity respectively and r_{igt}^{PE} is the individual specific return on private businesses. In this case time effects and the two portfolio shares will not exhaust variation in returns, which now have an individual specific component. In equation (7) we will thus control for the wealth share in listed stocks and the wealth share in private businesses separately. To capture the correlation between returns and wealth documented in Section (4) we add to the specification a full set of dummies for the individual wealth percentiles computed using lagged wealth values.

While the role played by observable characteristics is important, the residual term u_{igt} is our focus. We model the residual u_{igt} as being the sum of an individual fixed effect and an idiosyncratic component, which may possibly exhibit serial correlation. Hence:

$$u_{igt} = f_{ig} + e_{igt}$$

The fixed effects f_{ig} capture persistent differences in average returns across people belonging to a given generation g. This may arise from systematic differences in risk preferences, triggering different portfolio composition and thus returns on wealth, or from differences in ability or in opportunities to access investment alternatives - including systematic persistent differences in private businesses productivity. Because we observe several generations in our data, we can also study intergenerational persistence in returns fixed heterogeneity by estimating:

$$f_{ig} = \rho f_{ig-1} + \eta_{ig}$$

Table 2 shows the results of the estimates. The dependent variable is the return on total wealth in year t (expressed in percentages). The first column shows the results from a pooled OLS regression, without the fixed effects but adding a number of individual characteristics, some time invariant, to gain some intuition on the role played by the covariates of the returns. Observable heterogeneity in wealth returns is captured by demographics (gender, municipality fixed effects, age dummies, number of years of education, a dummy for economics or business education, employment, and marital status dummies), year fixed effects (to capture aggregate variation in returns), and the lagged shares in listed risky assets and in private business out of total wealth. We run these regressions on our total sample, comprising almost 50 million observations. The estimates show that households headed by a male have – ceteris paribus – a lower average return on wealth, but the effect is economically negligible (3.1 basis points). Returns are correlated with general education and with specific education in economics or business. A 10-year increase in general education results in 38 basis points higher returns on wealth and having received economics or business education is associated with 12.2 basis points higher returns. A systematic difference in returns of 38 basis points can produce a difference in wealth at retirement of 16.4% over a working life of 40 years. This effect is above any effect that education may have on portfolio returns because it twists the portfolio allocation towards riskier and more remunerative assets. This finding is consistent with Bianchi (2015) and von Gaudecker (2015), who find a positive effect of a measure of financial literacy on the return to investments among French and Dutch investors respectively, but with references to a specific asset. It also supports the results of Lusardi and Mitchell (2015) who study the effect of financial knowledge on returns to wealth and assets at retirement but within a life cycle model calibrated on US data.

Not surprisingly, portfolio shares in listed stocks and in private equity have both a positive and large effect on the return on wealth with the effect of the share invested in private equity significantly larger than the effect of the share in listed shocks. Increasing the share in listed stocks by 30 percentage points (about the move from the risky share of a non-participant to that of the average participant) increases the return on wealth by 20 basis points; increasing the share in private equity by the same amount is associated with a much larger increase in returns on wealth of 194 basis points. This finding is consistent with the idea that because private equity is highly concentrated, it yields a private equity premium. It runs contrary to that of Moskowitz and Vissing-Jorgensen (2002) who found a "private equity premium puzzle" (too low returns given the risks involved) using data from the US SCF; but is consistent with the results of Kartashova (2014) who documents the existence of a private equity premium using the same survey but extending the sample to the more recent waves. Overall, these estimates suggest that part of the observable heterogeneity in returns reflects compensation for risk for investing in listed stocks or for idiosyncratic risk in private businesses. Time fixed effects, though not shown, are always significant, as are age dummies. Interestingly, returns on wealth show an increasing pattern over the life cycle up until about age 60, and then a declining pattern in the years following retirement. Because we control for the share of wealth held in risky assets, this pattern does not reflect changes in portfolio composition over the life cycle. Yet, overall observable characteristics explain only 7% of the variance of individual returns to wealth. This limited fit is remarkable because, as noticed, the canonical portfolio model with fully diversified risky portfolios would imply that, controlling for time variation in returns, all heterogeneity in returns should be explained differences in the risky shares.

The second column modifies the specification by replacing the risky shares with their interaction with the time dummies. This more flexible specification captures differential effects of the risky share on individual returns as the aggregate component of returns varies. In addition the interaction between the share in private equity and the time dummies captures variation in individual returns due to tax-induced changes in incentives to distribute corporate dividends following the 2006 tax reform.[11]

The fit of the model improves (the R^2 increases from 0.077 to 0.115) but the size and significance of the other effects are otherwise unchanged. The third column adds the individual fixed effects. Because the model includes age and time effects, the individual fixed effects also capture cohort effects, posing a well known identification

[11]In Norway until 2005 distributed dividends were essentially exempt from tax (expect for an 11% tax in 2001) while capital gains were taxed at the same rate of 28% as retained profits. A reform passed in 2006, but anticipated since at least 2001, has moved to a new regime in which distributed dividends are taxed at the same marginal tax rate as earned income, at least for the part of returns on equity exceeding a risk free return of 3%. The corporate tax rate was kept at 28% (see, Alstadsæter et al. (2006)). The interaction between the time effects and the share of wealth in private equity captures the fact that private business owners may have timed the distribution of dividends in response to changes in tax incentives.

problem arising from the linear relation between age, time and year of birth. We deal with this issue by using the Deaton and Paxson (1994) restriction and impose that time effects sum to zero once the variables have been detrended. Since our data cover several years, we are able to separate trend and cycle, and thus are reasonably confident about the decomposition of age, time and cohort effect based on this restriction (Deaton 1997). As usual, the effect of time-invariant characteristics (such as gender or education) is no longer identified. The portfolio share in risky listed and non-listed assets can be identified out of time variation. The effect of of the share in listed stocks is now larger and that on private equity smaller: the effect of a 30 percentage points increase in the share of listed stocks results in a 30 basis point increase in the return on wealth and an equal increase in the share in private business is associated with an increase in returns of 137 basis points. The key feature of this regression is that the individual fixed effects improve the fit substantially by tripling the R^2 compared to the regression in column 1, implying that returns have an important persistent individual component. The last column uses the specification in column 2 by allowing for interactions between the time effects and the risky shares. With this flexible specification and the individual fixed effects the model can explain more than a quarter of the variance of individual returns to wealth.

From $u_{igt} = f_{ig} + e_{igt}$, additional persistence in returns may in principle come from e_{igt}. To check whether this is the case, we look at the evolution of moments of the residuals in first difference, i.e., $E(\Delta u_{igt} \Delta u_{igt-s})$ for $s \geq 0$ (since taking first differences of the residuals removes the fixed effect, i.e., $\Delta u_{igt} = \Delta e_{igt}$). We find that these moments are minuscule and economically undistinguishable from zero for $s \geq 2$, consistent withe_{igt} being serially uncorrelated (see the autocovariances plotted in Figure A.1 for all values of s).

To highlight the role played by observed and unobserved heterogeneity in explaining individual returns to wealth we use the estimates in Table 2, column 3 to compute the components of $E(r_{igt}|P_w) = E(X_{igt}|P_w) + E(f_{ig}|P_w) + E(e_{igt}|P_w)$ where P_w denotes the wealth percentile; averages are taken across individuals and over time. Figure 13 shows the elements of the decomposition. Average returns are, as already documented, increasing in wealth, and at an increasing speed at the top of the distribution. Observed heterogeneity explains a large part of average returns but its contribution is flat over a wide range of the wealth distribution. However, it contributes considerably to explain the large increase in returns among the highest wealth percentiles. This is largely the reflection of the share in private equity which, as shown in Figure, dominates the portfolio of people at the very top.

The fixed effect component is correlated with wealth over the whole range of the wealth distribution, and is thus the main driver of the positive correlation between returns and wealth.

6 Heterogeneity and persistence in returns to wealth

6.1 Within generation heterogeneity and persistence

Figure 14 plots the empirical distribution of the individual fixed effects (obtained from the estimates in the third column of Table 2). Because we include a constant in the regressions, the distribution is centered around zero. It has a long right tail[12] and is quite disperse, with a standard deviation of 2 percentage points and a 90th-10th percentile difference of 4 percentage points.

One interesting question is whether the persistent component of the wealth returns is associated with observable characteristics that, a priori, one may consider economically relevant. In Figure 15 we thus plot the distribution of fixed effect for business owners and non-owners (first panel); top vs. bottom wealth groups (second panel); more vs. less educated individuals (third panel), and for people with and without an economics or business degree (last panel). Because the first two characteristics (being a business owner and being at the top of the wealth distribution) may vary over time, the non-owners and those in bottom wealth groups are defined using indicators for "never being a business owner " and "never being in the top 10% of the distribution". Independently of the grouping considered, there is substantial heterogeneity in estimated fixed effects. Group differences are also economically significant. Business owners exhibit a distribution of persistent returns that is much more spread out and shifted to the right, which is consistent with owners of private equity facing more heterogeneous investment opportunities and higher returns on capital. Returns are heterogeneous both among the wealthy as well as among people at the bottom of the wealth distribution. But the distribution of persistent returns in more spread out and returns are on average higher among the wealthy, with differences in mean and spread becoming larger at the very top of the wealth distribution. Individuals with more general schooling have a much less dispersed distribution of persistent returns to wealth while those wth a degree in economics or business face both more dispersion in persistent returns with a distribution more shifted to the right, consistent with a positive correlation with education.

[12]For visual clarity we collapse the frequency mass of fixed effects above **12 percentage points and below XXX**.

6.2 Intergenerational persistence

Benhabib et al. (2011, 2015) argue that for return heterogeneity to be able to explain empirically the thickness in the right tail of the stationary distribution of wealth, individual average returns on wealth need to display some persistence across generations. Our empirical strategy allows us to calculate the fixed effects of returns on wealth for the whole adult population of Norway. Since the sample includes members of different generations (parents and children, etc.) at least for some of the years, we obtain measures of the wealth returns fixed effects for both parents and children whenever they belong to different households.[13] This allows us to test whether wealth returns are correlated across generations, and to check whether such correlation is coming from the persistent component or from observable characteristics that may be shared by both generations.

We start by ranking parents according to their financial wealth, the return to it, and the persistent component of the returns (fixed effect). For each variable, we allocate parents to various wealth percentiles. We do this by cohort (year of birth) and year. Next, for each percentile of the parents' variable of interest (wealth, returns, or return fixed effect), we compute the average percentile occupied by their child in the distribution of the same variables in the same year (again, relative to their year of birth cohort to control for the fact that parents and children are in different points of their life cycles). The three panels in Figure 16 show the correlation between the wealth percentile of the parents and that of the sons (top left panel), between the returns percentiles (top right panel), and between the permanent components of these returns (bottom panel).[14]

The figure reproduces the well-known positive association between parents' and sons' wealth ; as in other studies, intergenerational correlation in wealth, while high, is far below the perfect intergenerational correlation benchmark (implying no wealth mobility) of the 45-degree relationship marked by the green solid line (corresponding to a regression coefficient of 1). Indeed, a regression of the father's rank on the average rank of the son has a coefficient of 0.29 (s.e. 0.006) (the dashed line in the graph). There are also some clear non-linearities, which the simple OLS regression ignores. Intergenerational persistence in wealth is much stronger at the top of the father's distribution of wealth, a feature that can be detected precisely in our data because of their universal coverage. Some of the low correlation may be partly due to the fact that we are measuring variables when parents and children are both alive

[13] As mentioned above, we focus on fathers and sons.

[14] We are able to retrieve the fixed effects for almost 2 million (1,959,956) parents and their corresponding children.

(i.e., before inheritances).

As shown in the second panel, the average percentile of the son's return to wealth is also positively correlated with the percentile return of the father. But this correlation is much milder than that for overall wealth. While a father at the 80th percentile of the wealth distribution predicts a son at the 60th percentile, in the case of returns there is greater reversion to the mean: a father at the 80th percentile of the distribution of returns predicts a son only at the 54th percentile on average. A linear regression of the father's rank onto the average child's rank has a coefficient of 0.08 with a standard error of 0.003. Here too, there are important non-linearities. First, the bottom of the distribution is perhaps affected by measurement error. Second, the relation turns negative at the very top of the parents' returns. Intergenerational correlation exists also in the persistent component of wealth returns. The third panel shows that the intergenerational correlation in returns is very similar when we use the persistent component of the son's return to wealth (the fixed effect) against the percentile of the persistent component of the father's return. The linear regression has a coefficient of 0.09 with a standard error of 0.003. Sons of extraordinary parents in terms of returns to wealth over their life cycle tend to very quickly revert back to the mean.

Obviously, given the positive correlation between returns and wealth, all or part of the intergenerational correlation in returns documented in the figures may just reflect intergenerational correlation in wealth. Alternatively, it may reflect the fact that since parents and sons overlap for several years they may be subject to the same aggregate shocks to returns. The positive correlation in the third panel of Figure 16, between sons' and parents' return fixed effects, rules outs the second possibility but not the first. To deal with this we report controlled regressions of sons' returns on fathers' returns. We show the results in Table 3; the first panel uses sons' and fathers' return percentiles as the left hand side; the second panel shows returns directly. The first column has no controls; as reported in Figure 12, the slope coefficient is small. All the other regressions include wealth percentile dummies **.** Adding wealth controls and age dummies lowers the slope of the intergenerational relation. However, it remains positive and significant. Results are unaffected when individual controls are added (third column). Including individual fixed effects in the last column flattens farther the relation but raises considerably the fit (the R^2 increases from 0.062 to 0.37 in panel A and from 0.05 to 0.25 in panel B) suggesting that the intergenerational correlation in returns is driven by the permanent component of returns. Table 4 shows the transition matrix when we allocate individuals (fathers and sons) according to their returns fixed effect in quintiles (relative to their year of birth cohort). There is similar persistence across different parts of the distribution. A son born to a parent in the top quintile has a 24 percent probability of also being in the top quintile (relative to individuals of his age), and an 18 percent probability of slipping to the bottom quintile. Overall, our data suggest substantial persistence and heterogeneity in returns within a generation but smaller persistence across generations. This result is similar to that found by Benhabib et al. (2016). In their exercise there is little evidence of intergenerational persistence in returns. In our case we have much more statistical power and indeed we find an economically small but statistical significant degree of persistence.

7 Assortative mating in returns to wealth

In the literature there is evidence of assortative mating by education, income, and parents wealth (Eika et al. (2014); Greenwood et al. (2016); Lam, 1988; Charles et al., 2013). As far as we know, there is no evidence about assortative mating on personal wealth or returns to wealth. Both can be studied with our data. Differently from assortative mating on stable characteristics such as education or even parents' wealth (which do not vary because of marriage and can thus be measured *after* people have mated), detecting assortative mating on personal wealth and returns to wealth requires that these variables are observed *before* marriage, which explains the practical absence of any empirical evidence on it. After marriage, individual wealth and returns to wealth are hard if not impossible to separate. The large sample size and the long time span of our data allow us to identify and follow individuals for several years before they get married and thus test whether individual wealth and pre-marriage returns of future spouses or partners are correlated as implied by assortative mating. As mentioned in Section X.X, some of the assortative mating in returns to wealth may reflect the fact that individuals also sort on wealth.

To detect assortative mating on returns over and above assortative mating in wealth or other traits, we estimate the following model:

$$r^h_{it-pre} = \lambda r^w_{it-pre} + Z'_{it-pre}\mu + \varphi_{it-pre} \tag{8}$$

where r^j_{it-pre} is the average return to wealth for spouse j= $\{h, w\}$ measured in the years before marriage or shared custody of a child (excluding the two years before we firstly observe the individuals sharing a tax ID), Z_{it-pre} a vector of controls and φ_{it-pre}an i.i.d. error term.

Table 5 shows the results of the estimates. Before delving on assortative mating on returns to wealth, we document extensive assortative mating on personal wealth (see column 1). Pre-marriage wealth of the two future spouses appears to be strongly correlated with a slope coefficient of 0.24**.** This results is distinct from Charles et al (2013) who document that spouses sort on the basis of parental wealth. The other columns show results for the returns to wealth. The second column shows a positive correlation between the pre-marriage returns to wealth of the two spouses, with a slope of 0.11. This is consistent with people choosing to marry with a partner that has a similar ability to generate returns out of wealth. In column (3) we control for the possibility that the association between returns is spuriously due to assortative mating on other traits, such as education, aggregate effects (which we control for using year of marriage dummies), or life cycle stage (which we control for using age at marriage). The effect declines in magnitude but not in significance. In light of the positive correlation between returns and wealth, the positive correlation between the returns of the perspective spouses could reflect the assortative mating based on wealth documented in the first column. The other columns of Table 6 add increasingly finer controls to make sure assortative matching in returns is not just a spurious reflection of assortative mating on wealth. In column (4) we classify each spouse by whether they have above or below median pre-marriage wealth and then add dummies for whether both are above median ("both rich"), both below median ("both poor"), and so forth (with "both poor" being the excluded category). These dummies capture assortative mating in wealth. Controlling for them lowers slightly the slope of the mating relation in returns (from 0.08 to 0.07), but significance remains high. The remaining two columns add much finer controls for wealth mating: 25 dummies corresponding to the pairings of five wealth quintiles, or 10,000 dummies corresponding to the pairings when spouses are classified according to their wealth percentiles. Results remain similar, implying that assortative mating on returns is not just a reflection of assortative mating in wealth. Even whithin a narrow wealth group (say, husband and wife both in the top percentile), the husband's pre-marriage return to wealth is positively associated with the wife's pre-marriage return.

Assortative mating can be predicated on two possibilities. First, it may reflect yet another trait people match on – such as preferences for risk, entrepreneurial spirit, etc.. Second, it may be justified by a demand for wealth preservation: the wealth of a high-return individual may be threatened by the poor return of his/her spouse. For this to be the case the low-return spouse must contribute to the wealth management of the households once assets are jointly owned after marriage. If instead, following marriage, the management of household wealth is taken care of by the high-return spouse there is no scope for assortative mating. To test whether assortative mating on returns is justified, we regress the post-marriage (household) return against the pre-marriage returns of the two spouses. In particular, we consider the regression:

$$r_{it-post} = \beta_0 + \beta_1 min\{r^h_{it-pre}, r^w_{it-pre}\} + \beta_2 max\{r^h_{it-pre}, r^w_{it-pre}\} + Z^{'}_{it}\theta + \nu_{it}$$

where $r_{it-post}$ denotes the post-marriage household return, r^j_{it-pre} is the pre-marriage returns of spouse $j = h, w$, Z is a vector of controls and ν is an error term. Results of the estimates of this model are shown in Table 7. Interestingly, both the lower and higher pre-marriage return contribute to predicting the post-marriage household return. Their effect is hardly affected by the controls. The return of the spouse with the higher pre-marriage return has the strongest effect with a coefficient that is more than five times that of the return on the spouse with the lower pre-marriage return. But the latter matters, leaving room for some demand for assortative mating. Letting r^L =$min\{r^h_{it-pre}, r^w_{it-pre}\}$ and r^H =$max\{r^h_{it-pre}, r^w_{it-pre}\}$, we can summarize the effect of pre-marriage returns on post-marriage returns as $r_{it-post} = (\delta r^L + (1-\delta)r^H)$, where $\delta = \frac{\beta_1}{\beta_1+\beta_2}$ is the weight assigned to the spouse with the loweer return. Using the estimates in column (4), the weigths are 0.85 for the highest return and 0.15 for the lowest.

The last column adds interaction terms between the maximum and minimum return with a dummy for whether the highest return is that of the male, to accont for potential differences across genders in the weigth of the lowest and highest pre-marriage return. The results show that when the spouse with the higher pre-marriage return is the male, the effect of the higher pre-marriage return become higher (0.21) and that of the lower pre-marriage return is an order of magnitude smaller (0.02). The weights computed above are now $\delta = 0.09$ and $(1-\delta) = 0.91$. When instead the spouse with the highest pre-marriage return is the female, the two effects are respectively 0.16 and 0.05. The weights on the lowest and highest return become $\delta = 0.24$ and $(1-\delta) = 0.76$. In other words, males carry a higher weight both when their are better at generating returns as well as when they are not.

8 Implications of returns heterogeneity for the wealth inequality debate

To be added

9 Conclusion

In this paper we have made extensive use of Norwegian administrative records to study the properties of returns to wealth. These data have the noteworthy feature that they allow to retrieve both income from capital and asset holdings from tax records, making it possible to compute measures of returns to wealth for the population of the Norwegian households. Because the data are available for 20 years we can identify the fixed effect component of returns and quantify their persistent heterogeneity within a generation as well as the correlation in returns across generations. These measures are the theoretical counterparts of the notions of heterogeneity used in theoretical models of the wealth distributions and in calibrated models of wealth inequality by Benhabib, Bisin and Lou (2015). We document that returns to wealth are heterogeneous in the cross section and tend to be positively correlated with wealth. Both features are time varying. Returns to wealth have an important individual fixed component, implying substantial persistence within a generation. Returns to wealth tend also to persist across generations but the intergenerational persistence is quite mildsmall. Overall, these findings lend qualitative support to models of wealth inequality that stress heterogeneity in returns to wealth as key drivers of the unequal distribution of assets and their extreme correlation at the top. Interestingly, even quantitatively, our measure of persistent heterogeneity is not far from those estimated by Benhabib, Bisin and Lou (2015) from a †model calibrated on US data and so is the degree of intergernerational persistence. But our data reveal also features that have so far been neglected in models that emphasizes returns heterogeneity. Returns tend to be moderately correlated with wealth – a feature that should amplify the effects of returns heterogeneity on wealth inequality. But also their heterogeneity is higher among the wealthy, a feature that should facilitate downward wealth mobility. We plan to assess the role of these properties for wealth inequality in future work.

Appendix

References

Calvet, L. E., J. Y. Campbell, and P. Sodini (2007, October). Down or Out: Assessing the Welfare Costs of Household Investment Mistakes. *Journal of*

Political Economy 115(5), 707–747.

Eika, L., M. Mogstad, and B. Zafar (2014). Educational assortative mating and household income inequality. *NBER working paper*.

Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2016). Technology and the changing family: A unified model of marriage, divorce, educational attainment and married female labor-force participation. *forthcoming, American Economic Journals: Macroeconomics*.

Piketty, T. and E. Saez (2003). Income Inequality in the United States, 1913–1998. *The Quarterly Journal of Economics 118*(1), 1–41.

Saez, E. and G. Zucman (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *forthcoming, The Quarterly Journal of Economics*.

Reference list to be completed

Figures

Figure 1. Portfolio Composition: by percentile

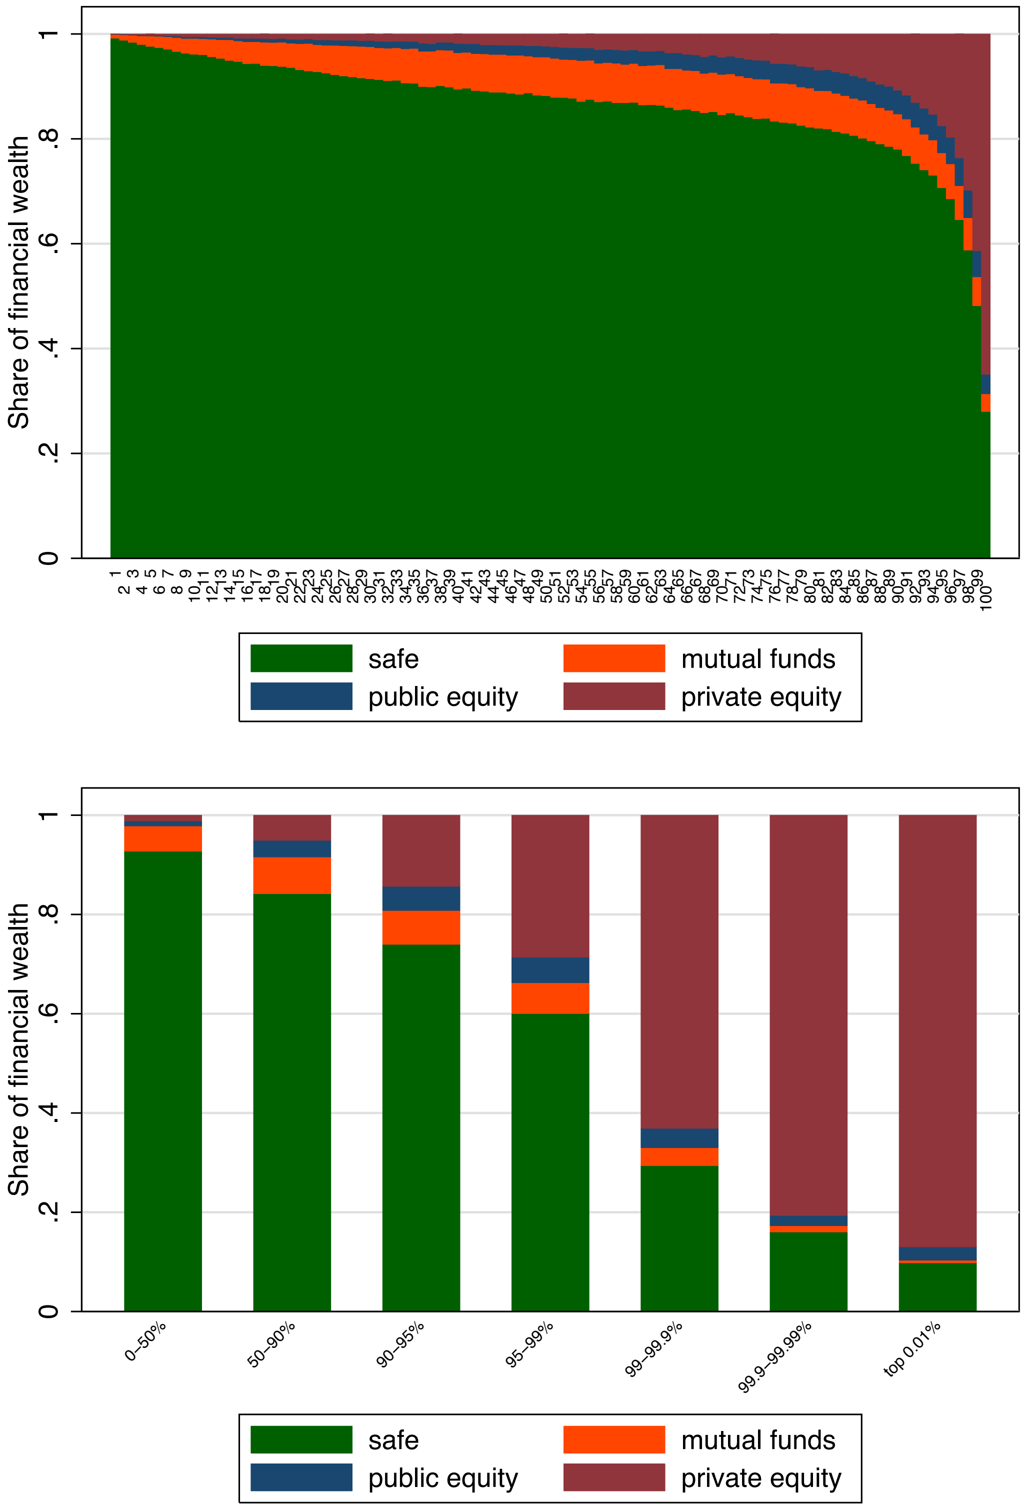

Notes:

Figure 2. Cross sectional distribution of returns

Notes:

Figure 3. Cross sectional standard deviation of returns by risky share

Notes:

Figure 4. Returns on Wealth over time

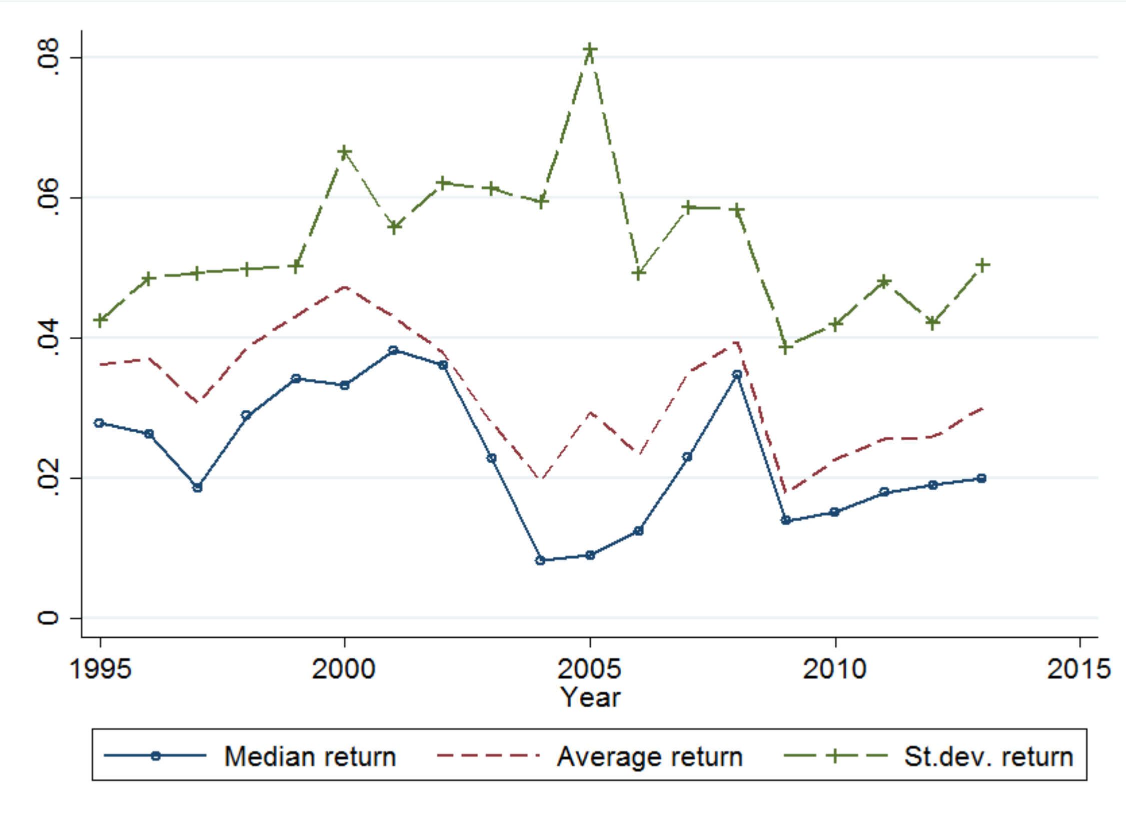

Notes:

Figure 5. Returns on Risky and Safe Assets over time

Notes:

Figure 6. Returns and Wealth percentiles (2013)

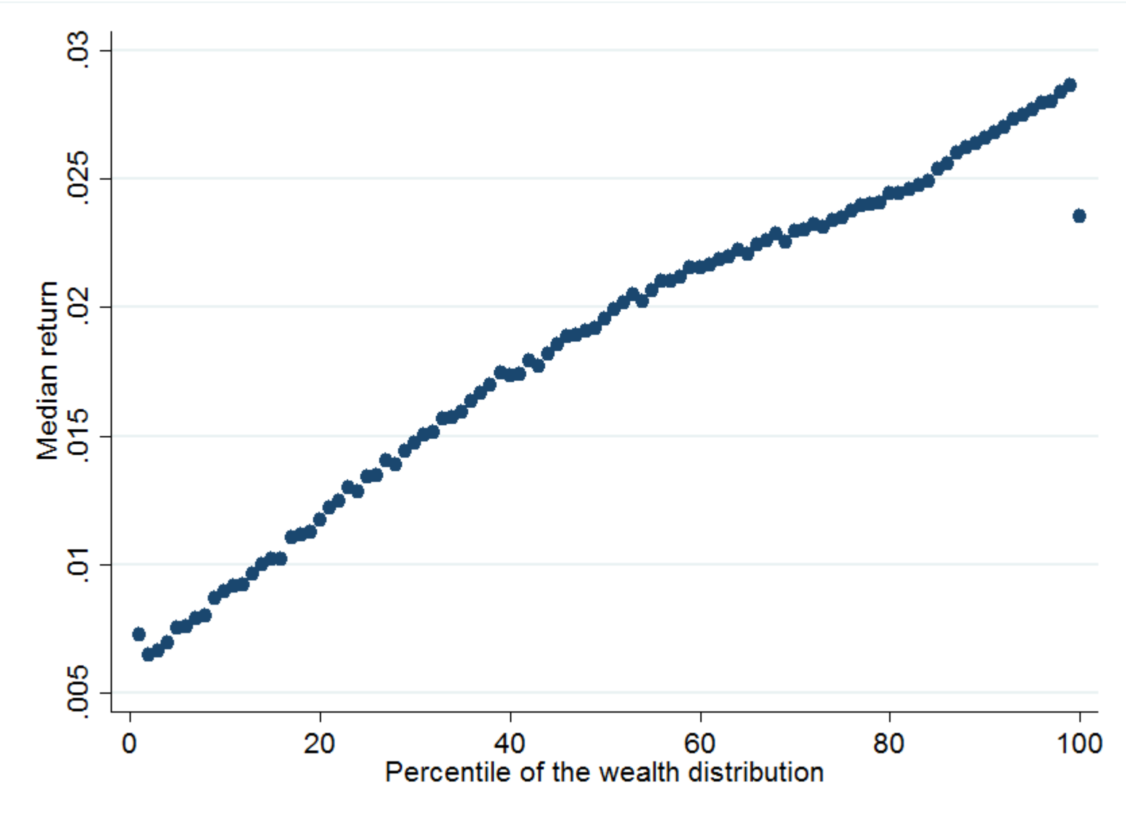

Notes:

Figure 7. Standard deviation of returns on wealth by wealth percentile in selected years

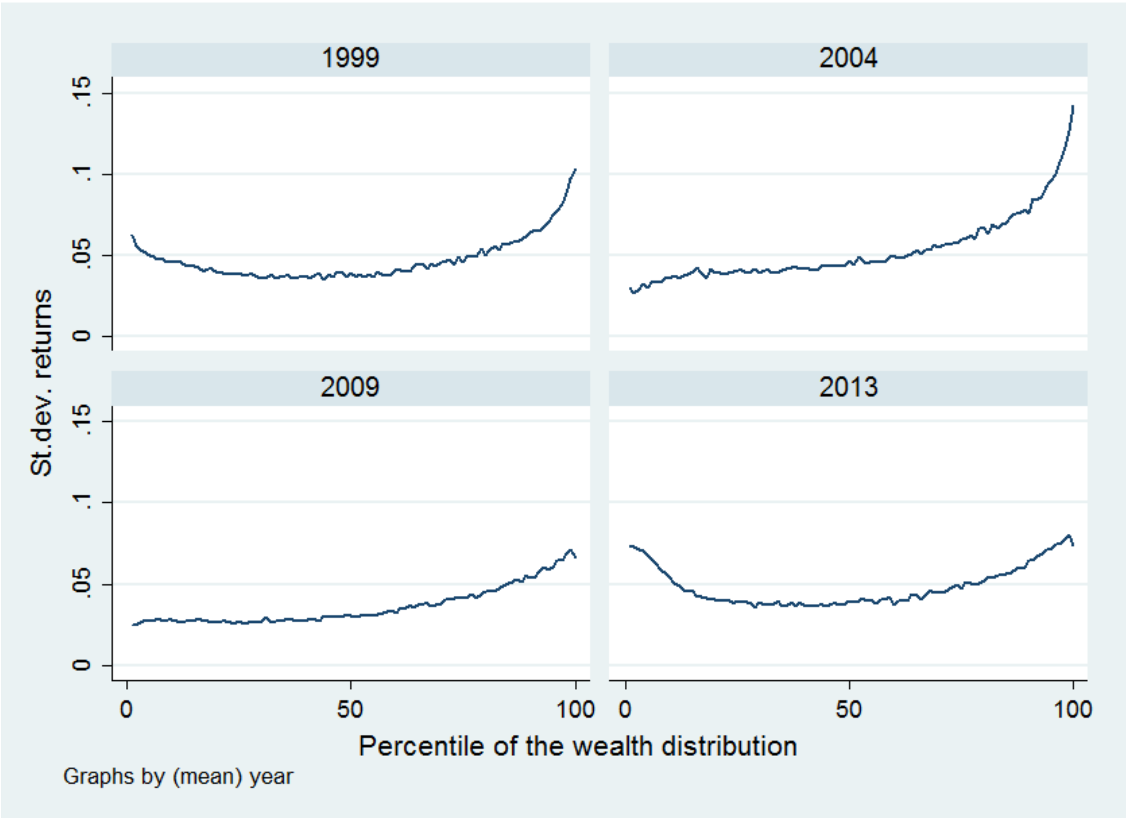

Notes:

Figure 8. Sharpe-Ratio vs wealth pctile 1995

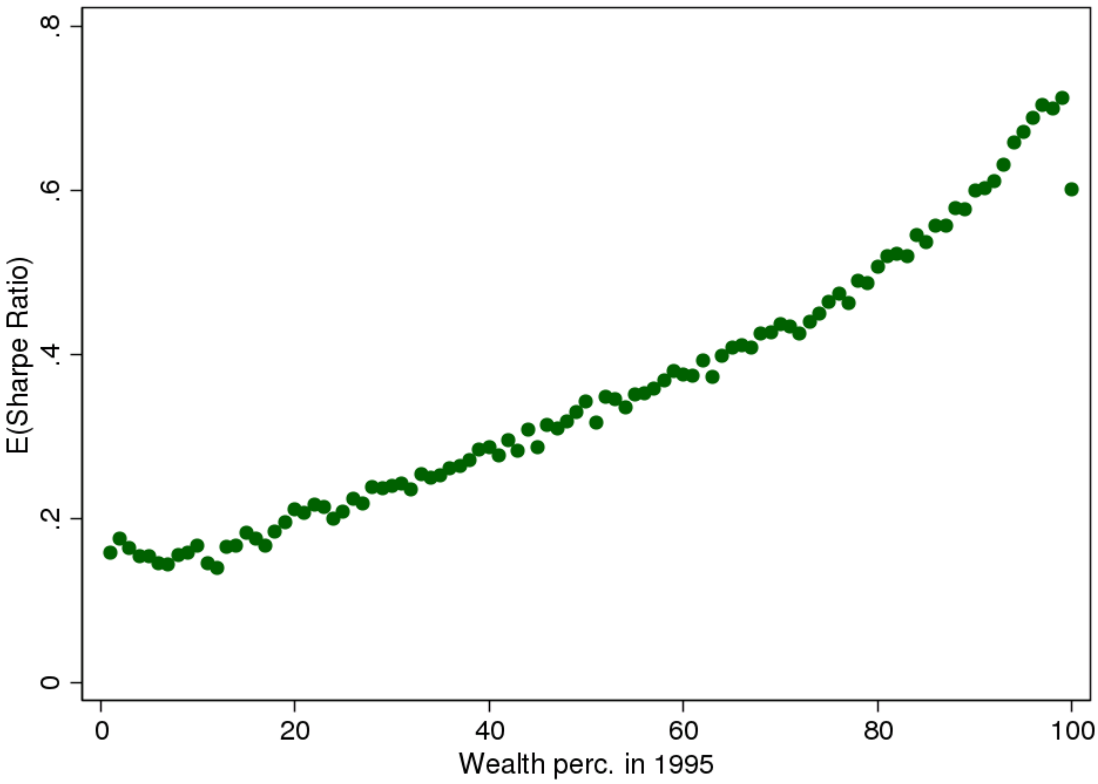

Notes:

Figure 9. Returns on Wealth over time by percentile in the wealth distribution

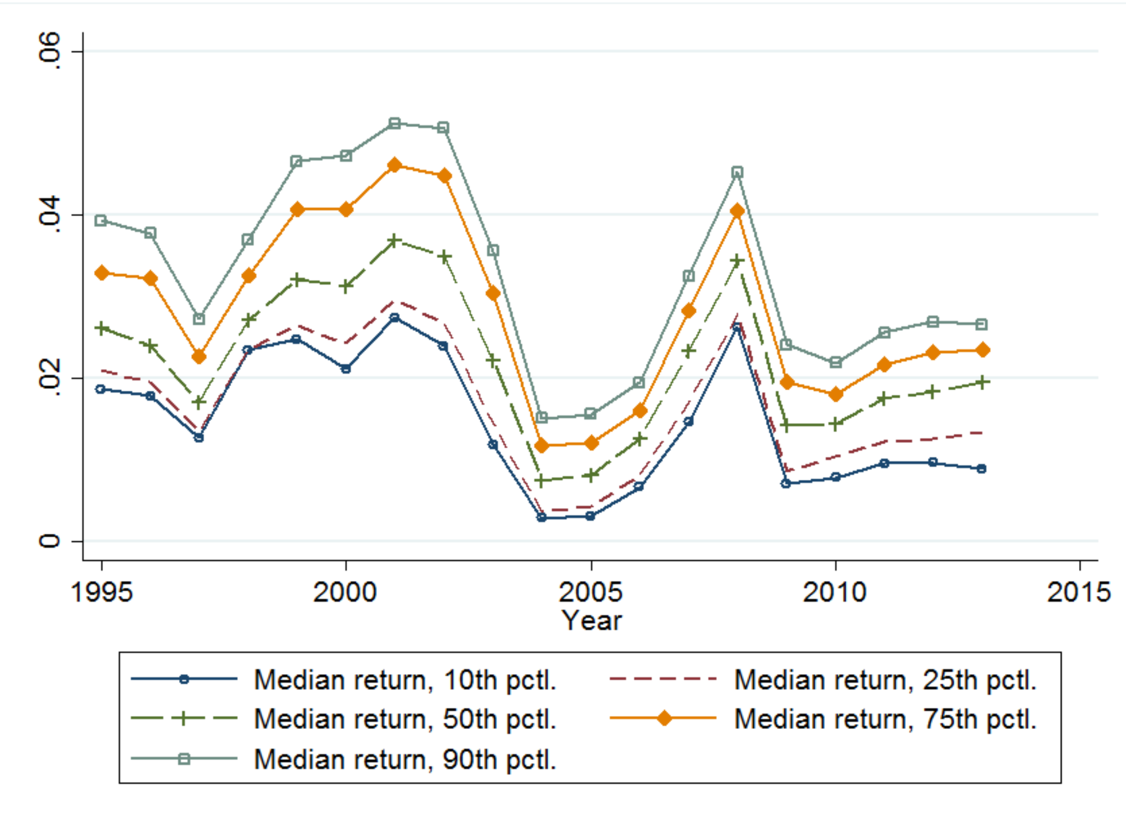

Notes:

Figure 10. Intergenerational correlation in returns percentile

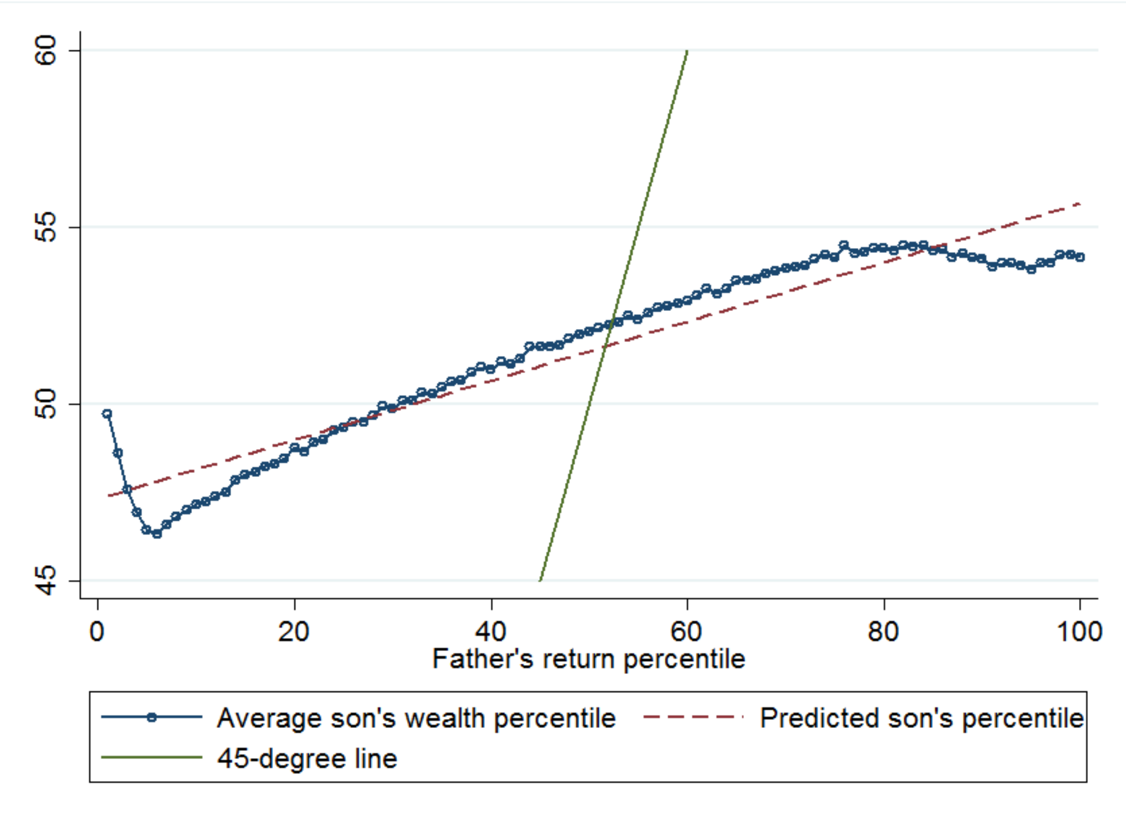

Notes:

Figure 11. Assortative mating on wealth

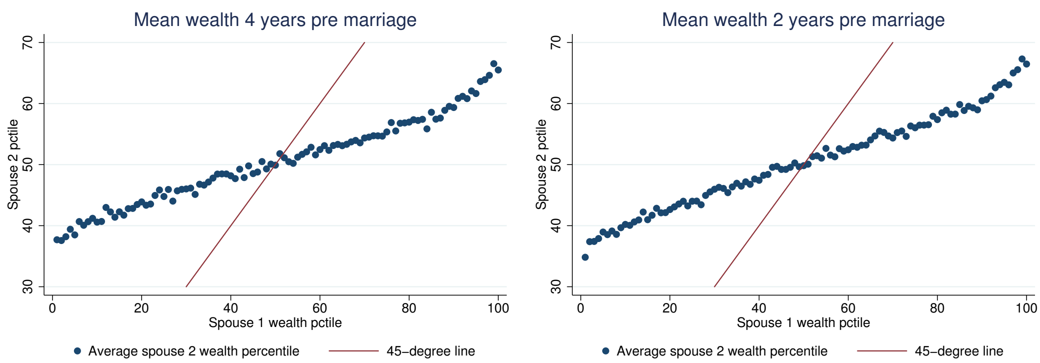

Notes:

Figure 12. Assortative mating on returns to wealth

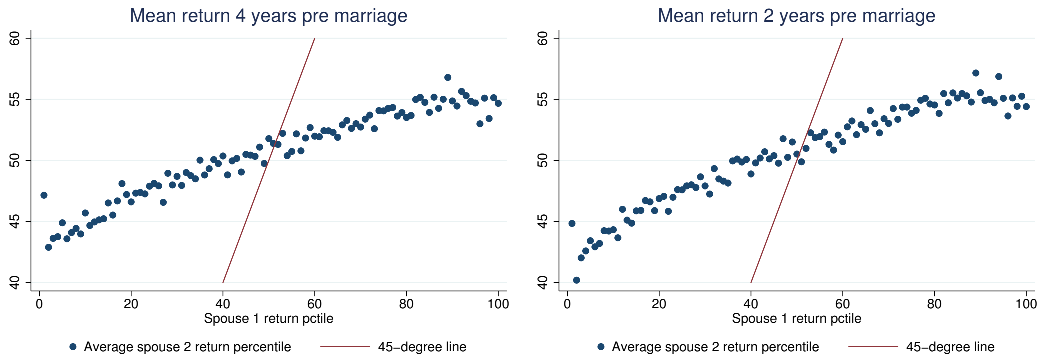

Notes:

Figure 13. Decomposing average returns by wealth percentile

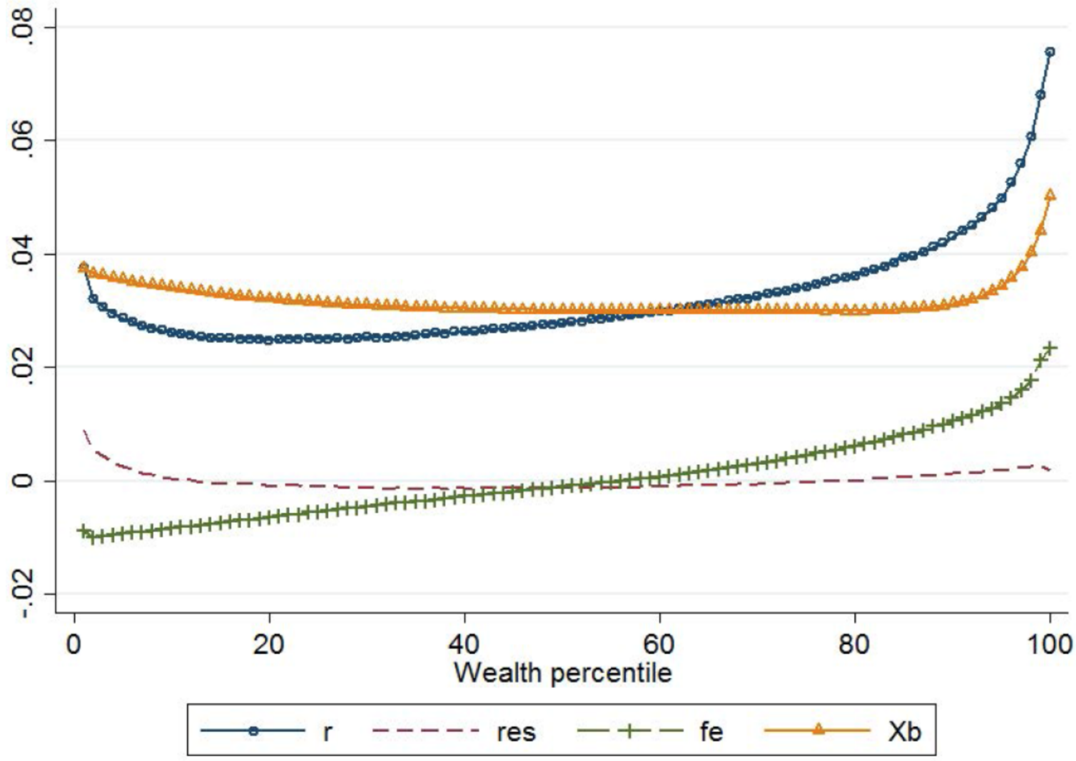

Notes:

Figure 14. Unconditional Distribution FE

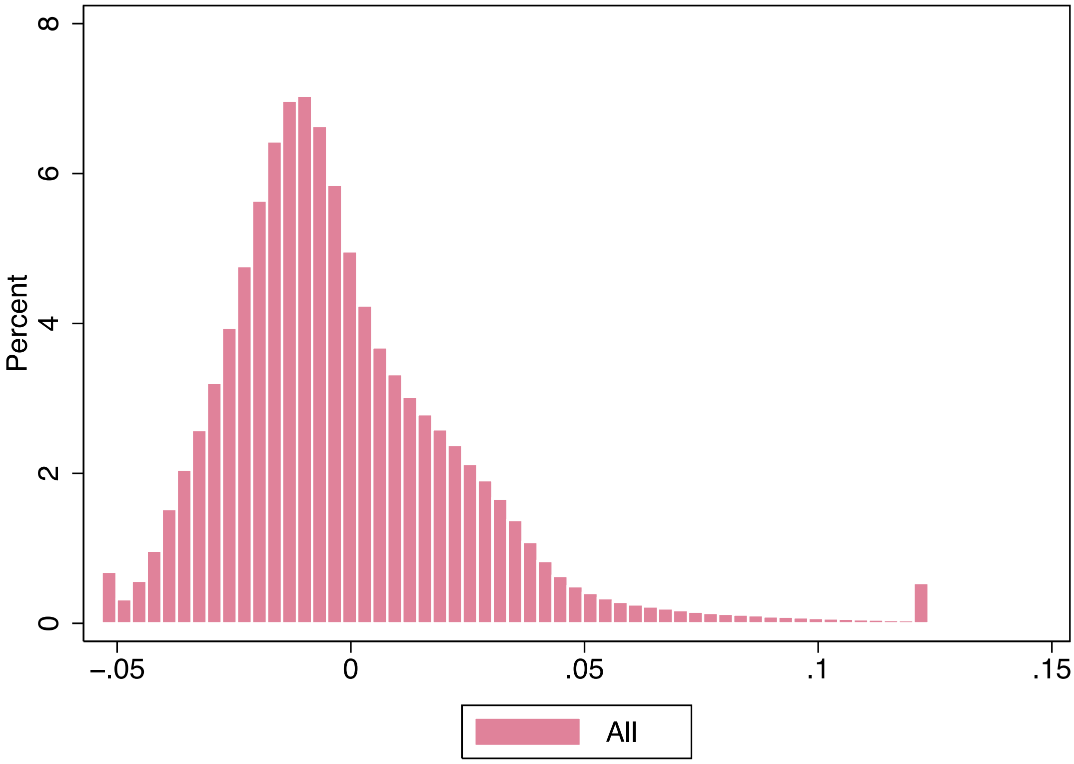

Notes: The figure plots....

Figure 15. Distribution FE broken down by observables

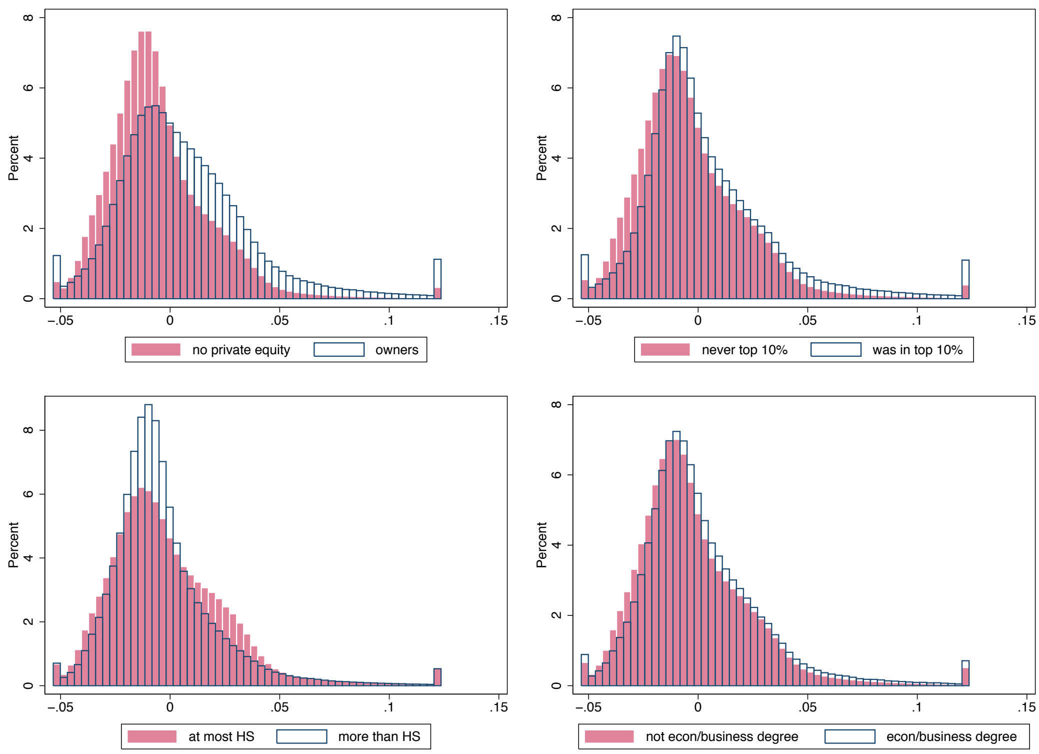

Notes: The figure plots....

Figure 16. Intergenerational correlation in FE percentile

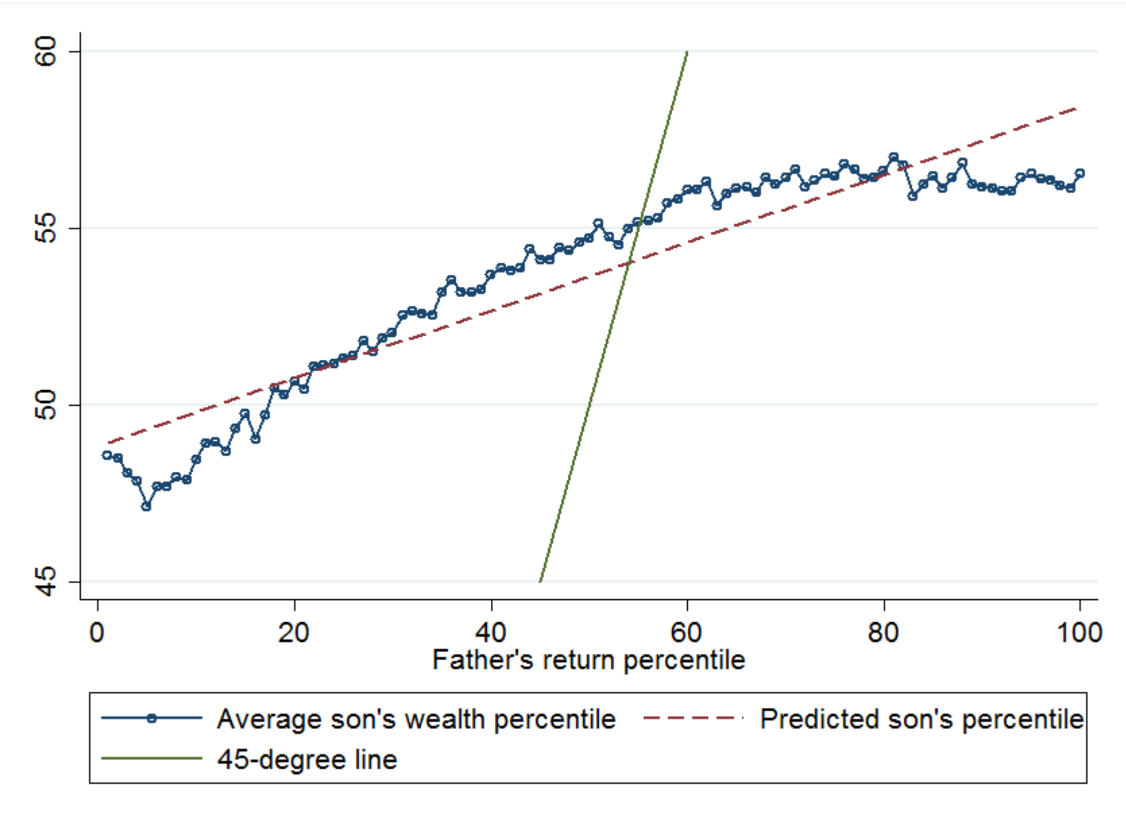

Notes:

Tables

Table 1. Summary statistics, 2013.

Panel A, Demographics:

	Mean	Std. dev	P10	Median	P90
Age	45.10	13.95	26	45	65
Male	0.50	0.50	0	1	1
Fraction married	0.49	0.50	0	0	1
Family size	2.70	1.34	1	2	5
Less than High School	0.19	0.39	0	0	1
High School	0.44	0.50	0	0	1
University	0.37	0.48	0	0	1
Years of education	13.74	3.64	10	13	17
Fraction with Econ/Business degree	0.12	0.32	0	0	1

Panel B, Assets and income:

	Mean	Std. dev	P10	Median	P90
Fraction w risky assets	0.45	0.50	0.00	0.00	1.00
Risky assets share	0.14	0.24	0.00	0.00	0.54
Cond. risky assets share	0.30	0.29	0.01	0.20	0.78
Fraction w private equity	0.11	0.32	0.00	0.00	1.00
Private equity share	0.05	0.18	0.00	0.00	0.05
Cond. private equity share	0.48	0.41	0.01	0.42	1.01
Fraction w public equity	0.38	0.49	0.00	0.00	1.00
Public equity share	0.09	0.20	0.00	0.00	0.35
Cond. public equity share	0.24	0.27	0.01	0.14	0.65
Risky assets	40,596	1,250,876	0.00	0.00	28,667
Safe assets	44,324	171,733	2,010	15,842	101,248
Total assets	84,920	1,321,363	2,291	19,957	142,451
Income from risky asset	1,979	46,784	0.00	0.00	421
Income from safe assets	1,161	5,115	10	320	2,683
Income from total assets	3,141	47,984	10	373	4,008

Panel C, Portfolio returns, percentages:

	Averages of portfolio returns, 2013		
	Total assets	Risky Assets	Safe Assets
Mean	0.030	0.058	0.025
Std. dev	0.050	0.234	0.032

	Averages of portfolio returns, 1994-2013		
	Total assets	Risky Assets	Safe Assets
Mean	0.032	0.035	0.030
Std. dev	0.054	0.256	0.033

	Value weighted averages of portfolio returns		
	Total assets	Risky Assets	Safe Assets
2013:	0.037	0.049	0.026
1994-2013:	0.050	0.071	0.032

Notes: Where applicable values in 2011 USD. Portfolio returns are reported in percentages. Averages of portfolio returns are calculated as the aritmetic means of the individual portfolio returns. Value weighted averages are calculated also taking into account the size of the individual portfolios. Public equities include stocks listed at the Oslo stock exchange and mutual funds.

Table 2. OLS and Fixed Effect regressions

	(1) Portfolio return b/se	(2) Portfolio return b/se	(3) Portfolio return b/se	(4) Portfolio return b/se
Lagged risky share	0.667***		1.014***	
	(0.009)		(0.012)	
Lagged private equity share	5.798***		3.554***	
	(0.023)		(0.024)	
Male	-0.031***	-0.031***		
	(0.002)	(0.002)		
Employed	0.224***	0.221***	0.011**	0.010**
	(0.003)	(0.003)	(0.005)	(0.004)
Years of education	0.038***	0.039***		
	(0.000)	(0.000)		
Econ/Business education	0.122***	0.121***		
	(0.004)	(0.004)		
Individual FE	no	no	yes	yes
Year FE	yes	yes	yes[1]	yes[1]
Age FE	yes	yes	yes	yes
County FE	yes	yes	yes	yes
Demographic controls	yes	yes	yes	yes
Lag. wealth percentile	yes	yes	yes	yes
Lag. risky share*i.year	no	yes	no	yes
Lag. private eq share*i.year	no	yes	no	yes
r2	0.077	0.115	0.233	0.267
N	47,686,405	47,686,405	47,686,405	47,686,405

Notes: Robust standard errors reported in parentheses. ***p<.01, **p<.05, *p<.10.

Table 3. Intergenerational returns regression.

Panel A: Returns (measured as percentiles):

	(1) q100ret b/se	(2) q100ret b/se	(3) q100ret b/se	(4) q100ret b/se
Father return percentile	0.082***	0.058***	0.055***	0.039***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	47.356***	47.029***	41.672***	54.537***
	(0.023)	(0.140)	(1.092)	(0.187)
Wealth controls	no	yes	yes	yes
Year FE	no	yes	yes	yes
Education length/type ind.	no	no	yes	no
Age	no	no	yes	yes
Individual FE	no	no	no	yes
r2	0.007	0.055	0.062	0.373
N	14,548,263	14,548,263	14,548,263	14,548,263

Panel B: Returns:

	(1) Portfolio return b/se	(2) Portfolio return b/se	(3) Portfolio return b/se	(4) Portfolio return b/se
Father's return	0.075***	0.050***	0.050***	0.046***
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	2.675***	3.388***	2.296***	3.087***
	(0.002)	(0.022)	(0.125)	(0.031)
Wealth controls	no	yes	yes	yes
Year FE	no	yes	yes	yes
Education length/type ind.	no	no	yes	no
Age	no	no	yes	yes
Individual FE	no	no	no	yes
r2	0.007	0.051	0.052	0.249
N	14,548,263	14,548,263	14,548,263	14,548,263

Notes: Standard errors reported in parentheses are clustered at the individual level. ***p<.01, **p<.05, *p<.10.

Table 4. Transition matrix

Notes: Standard errors reported in parentheses are clustered at the individual level. ***p<.01, **p<.05, *p<.10.

Table 5. Assortative mating on returns

	(1) Wealth pctile b/se	(2) Return pctile b/se	(3) Return pctile b/se	(4) Return pctile b/se	(5) Return pctile b/se	(6) Return pctile b/se
Wealth pctile spouse	0.240***					
	(0.002)					
Return pctile spouse		0.118***	0.083***	0.068***	0.067***	0.067***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Wealth Controls:						
Poor&Rich				-0.001		
				(0.002)		
Rich&Poor				0.145***		
				(0.002)		
Rich&Rich				0.151***		
				(0.002)		
Constant	0.380***	0.441***	0.206**	0.164*	0.136	0.244***
	(0.001)	(0.001)	(0.074)	(0.076)	(0.075)	(0.074)
Wealth controls		no	no	Rich x Poor	5 x 5	100 x 100
Age ind.		no	yes	yes	yes	yes
Marital year ind.		no	yes	yes	yes	yes
Education length/type ind.		no	yes	yes	yes	yes
r2	0.058	0.014	0.062	0.123	0.145	0.202
N	165,860	165,860	165,860	165,860	165,860	165,860

Notes: Robust standard errors reported in parentheses. ***p<.01, **p<.05, *p<.10.

Table 6. Assortative mating on returns, continued

	(1) ret__post__mar b/se	(2) ret__post__mar b/se	(3) ret__post__mar b/se	(4) ret__post__mar b/se	(5) ret__post__mar b/se
Pre-marital return of spouse with lower returns	0.020***	0.037***	0.036***	0.034***	0.046***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)
Pre-marital heturn of spouse with higher returns	0.190***	0.191***	0.190***	0.190***	0.155***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Age at marriage, female			-0.015***		
			(0.002)		
Age at marriage, male			0.011***		
			(0.002)		
Male*Return of higest					0.054***
					(0.005)
Male*Return of lowest					-0.025***
					(0.009)
Constant	2.274***	2.644***	2.753***	2.660***	2.695***
	(0.013)	(0.031)	(0.044)	(0.070)	(0.070)
Indicator controls:					
Marital year	no	yes	yes	yes	yes
Ages at marriage	no	no	no	yes	yes
r2	0.028	0.034	0.034	0.035	0.035
N	245,999	245,999	245,999	245,999	245,999

Notes: Robust standard errors reported in parentheses. ***p<.01, **p<.05, *p<.10.

Figure A.1. Unreported coefficients from Table 2: i.Year interaction with risky share and private equity share

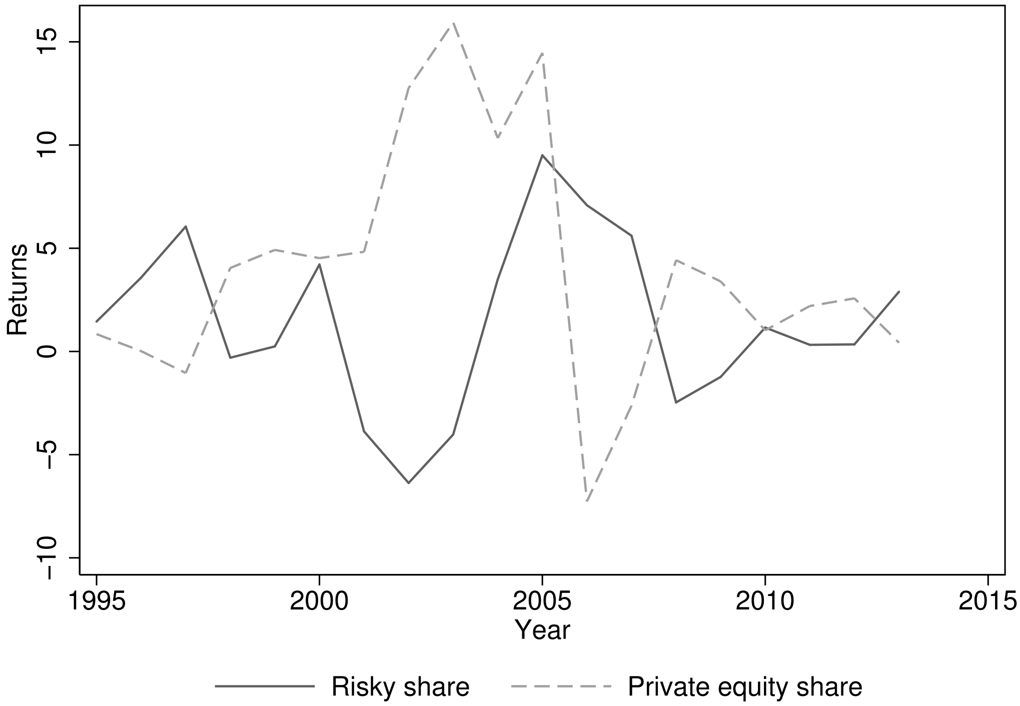

Notes: The figure plots....

Figure A.2. Auto-covariance first difference residuals FE regression

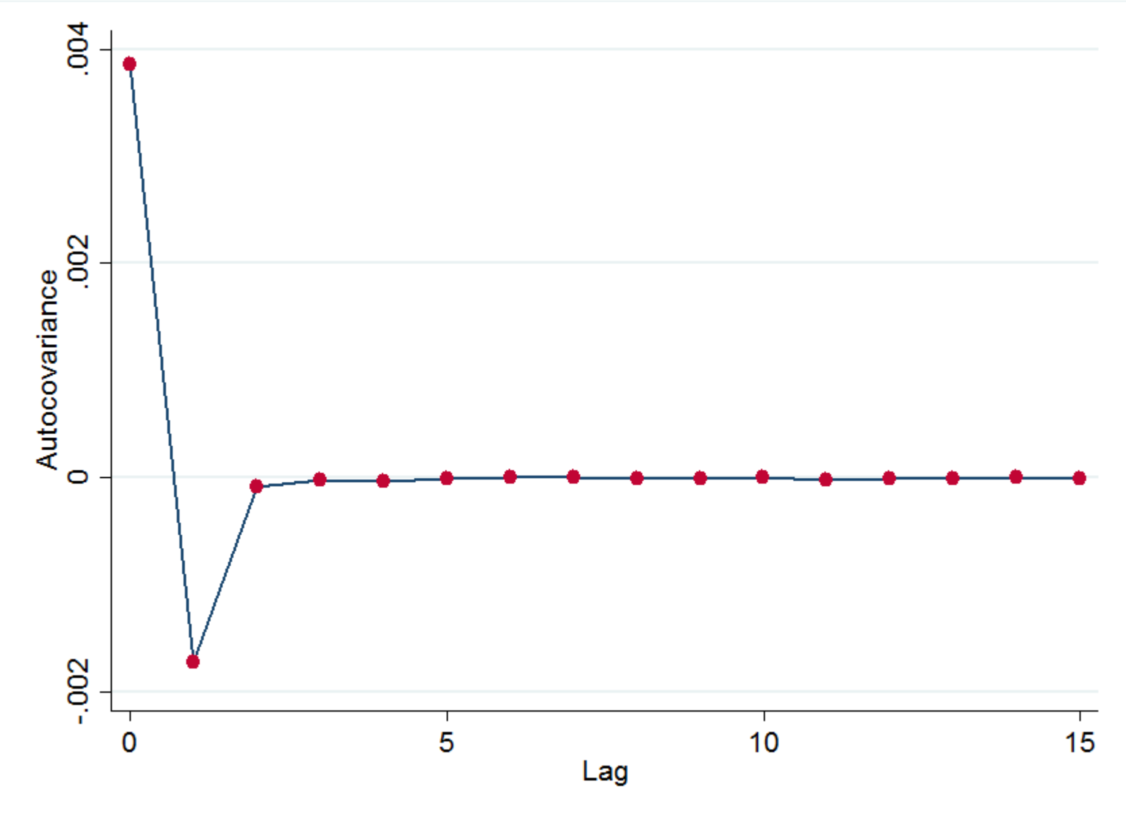

Notes: The figure plots....

Figure A.3. Concentration asset classes: by percentile

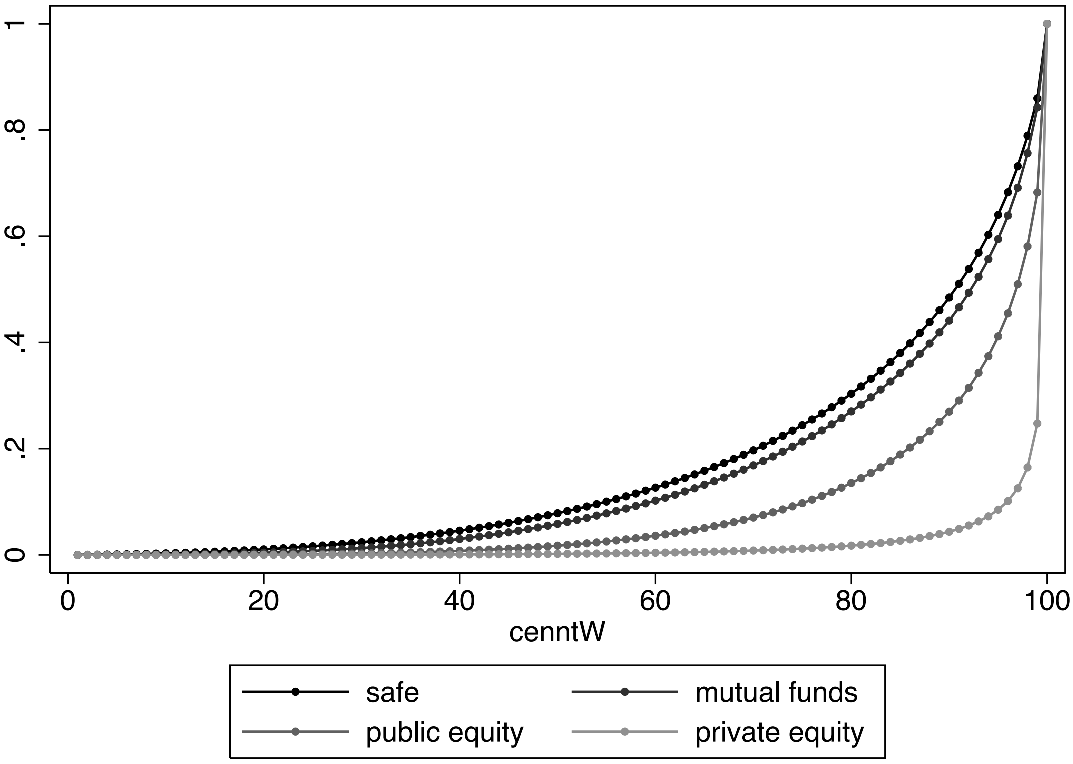

Notes: The figure plots....

Figure A.4. Concentration asset classes: by group

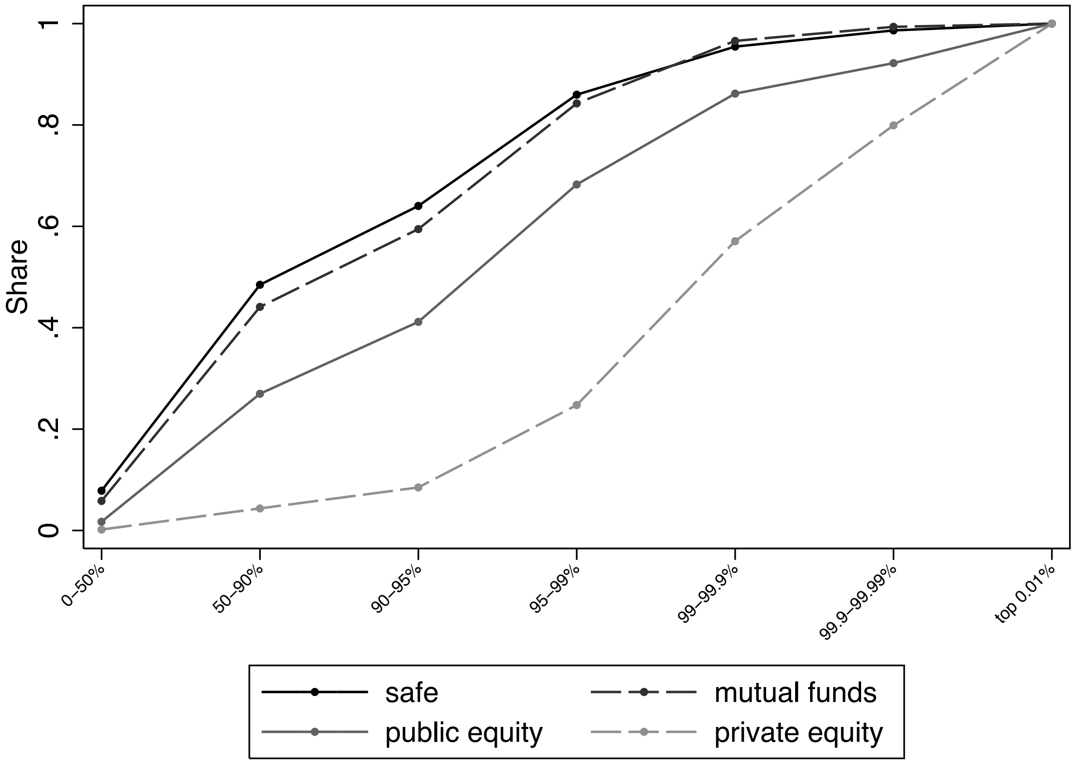

Notes: The figure plots....