The Life-Cycle Dynamics of Wealth Mobility*

Richard Audoly Rory McGee Sergio Ocampo Gonzalo Paz-Pardo

January 2025

Abstract

We use 25 years of tax records for the Norwegian population to study the mobility of wealth over people's lifetimes. We find considerable wealth mobility over the life cycle. To understand the underlying mobility patterns, we group individuals with similar wealth rank histories using agglomerative hierarchical clustering, a tool from statistical learning. The mobility patterns we elicit provide evidence of *segmented mobility*. Over 60 percent of the population remains at the top or bottom of the wealth distribution throughout their lives. Mobility is driven by the remaining 40 percent, who move only within the middle of the distribution. We show that the joint dynamics of income and wealth are inconsistent with standard models of savings. Instead, both persistent return and savings rates heterogeneity are necessary to account for wealth mobility, even for individuals away from the top of the wealth distribution. We show parental wealth is the key predictor of who is persistently rich or poor, while human capital is the main predictor of those who rise and fall through the middle of the distribution.

*Audoly: FRBNY; audolyr@gmail.com; Web: www.newyorkfed.org/research/economists/Audoly.
McGee: UWO & IFS; rmcgee4@uwo.ca; Web: https://sites.google.com/view/rorymcgee.
Ocampo: UWO; socampod@uwo.ca; Web: https://sites.google.com/site/sergiocampod.
Paz-Pardo: ECB; gonzalo.paz_pardo@ecb.europa.eu; Web: https://www.gonzalopazpardo.com.
The views below are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, the European Central Bank, or the Eurosystem.
We are grateful for comments and support from Roberto Iacono, Paolo Piacquadio, Alfred Løvgren, and the staff at the Oslo Fiscal Studies Centre at the University of Oslo, as well as from Viola Angelini, Javier Birchenal, Jim Davies, Mariacristina De Nardi, Jeppe Druedahl, Jan Eeckhout, Eric French, Michael Graber, Victoria Gregory, Fatih Guvenen, Amy Handlan, Juan Herreño, Joachim Hubmer, John Bailey Jones, Bariş Kaymak, Moritz Kuhn, David Lagakos, Hannes Malmberg, Elena Manresa, Cormac O'Dea, Serdar Ozkan, Tom Phelan, David Price, Pascual Restrepo, Baxter Robinson, Sergio Salgado, Lisa Tarquinio, and David Wiczer. We also thank seminar participants at various institutions and conferences. We also thank Emmanuel Murray Leclair for research assistance. Ocampo acknowledges financial support from the Research Council of Norway through the project TaxFair, Number 315765. McGee and Ocampo acknowledge financial support from the Social Science and Humanities Research Council of Canada through the Insight Development Grant Number 430-2022-00394.

1. Introduction

Do rich and poor people remain this way throughout their lives? Is it typical for people to experience reversals of fortune moving up or down the wealth distribution? If so, how large are the reversals, and when do they happen? These movements across the wealth distribution reflect the outcomes of critical events and choices in people's lives, including their human capital accumulation, earnings, and business activities. Wealth mobility thus speaks to the opportunities that people face.[1] However, despite growing evidence on the dynamics of wealth concentration for the wealthiest,[2] we know little about the life-cycle dynamics of wealth mobility for the population as a whole and the different income and savings patterns that generate this mobility.

Our main contribution lies in documenting wealth mobility over the life cycle. We conduct a comprehensive study of the complete distribution of lifetime individual wealth trajectories, which we construct using 25 years of administrative data from the Norwegian tax registry (1993–2017). We find *increasing wealth mobility* over the life cycle, so that an individual's initial position in the wealth distribution matters less as they age. Only one-quarter of individuals are in the same quintile of the wealth distribution after 25 years, whereas almost half of individuals stay in the same income quintile. However, this population trend does not, by itself, tell us much about the underlying life-cycle patterns that drive it. Who is actually moving? And how?

To answer these questions, we elicit typical life-cycle wealth trajectories from the distribution of wealth histories using agglomerative hierarchical clustering, a tool from statistical learning that groups individuals based on their entire realized trajectories.

[1] Low wealth mobility can be a symptom of limited equality of opportunity and can exacerbate the effects of high inequality. In the context of income inequality, Alan Krueger, then Chairman of the Council of Economic Advisors under President Obama, remarked that "*if we had a high degree of income mobility we would be less concerned about the degree of inequality in any given year*" (Krueger 2012, pg. 3).

[2] See Gomez (2023) for evidence from the *Forbes* 400 list and Ozkan, Hubmer, Salgado, and Halvorsen (2023) for evidence on the top 0.1 percent of Norwegian wealth holders. Quantitative analysis of the origins of the wealthiest individuals dates back at least to Wedgwood (1929).

With this, we are the first to flexibly and non-parametrically characterize wealth mobility over the life cycle. Our methodology makes it possible to characterize mobility beyond the transitions between any two periods in a way that captures much more information while still providing a tight and tractable characterization of the data. We cluster individuals into four main groups, whose typical trajectories explain more than one-half of the variation in wealth histories. We also study the heterogeneity within these groups by exploiting the hierarchical nature of our clustering methodology.

The mobility patterns we uncover show that increasing wealth mobility over the life cycle is not broad-based and is not driven by movements spanning the whole distribution. Instead, it comes from a combination of two largely immobile groups (60 percent of the population) that stay relatively rich and poor, and two groups that undergo large transitions that are nevertheless contained within the middle of the wealth distribution.[3] The two groups driving increasing mobility along the life cycle experience a reversal of fortunes as they age; one rising through the middle of the distribution and the other falling. These patterns have important implications for the majority of individuals in the population. While it is well known that the bottom 50 percent of individuals has zero to very little net worth, we are the first to document that more than 40 percent of the population is not only wealthless, but stays that way throughout their lives. We interpret these patterns as evidence of *segmented wealth mobility*: mobility takes place only for some groups of individuals and within a section of the distribution.

Further, we establish how different economic factors—such as portfolio composition, sources of income, family structure, and inheritances—relate to the large gaps in wealth accumulation between groups. We find that, while property is the primary asset for all groups, there are important differences in business assets and private equity. These

[3]Differences in realized wealth rank trajectories correspond to meaningful differences in levels reflecting the high degree of wealth inequality in Norway. For instance, the wealth gap at age 55 between the two groups in the middle of the distribution is 600,000 US dollars.

assets are concentrated in the top and falling groups, which aligns with their higher rates of self-employment. Notably, risers engage in less business activity and instead rely on employment income as they move up the distribution. Their labor income is higher than that of the fallers, and their household incomes match those of the top group (who have a larger share of capital income). By contrast, fallers have similar, but ultimately less successful, entrepreneurial activities compared to those at the top. The 40 percent of individuals at the bottom of the distribution are very different from the others: their incomes are persistently lower, they remain renters throughout their lives, and they rarely own businesses.

We also contribute by showing how the joint dynamics of wealth and income mobility add to our understanding of the drivers of wealth accumulation. We show that the inability of flexible income risk to generate levels of top wealth inequality extends to generating the saving behavior and wealth mobility of individuals across the entire wealth distribution.[4] Specifically, we show that standard life-cycle models of buffer-stock savings (Zeldes 1989; Deaton 1991; Carroll 1992) impose a tight link between income and wealth mobility, driven by differences in permanent income, which is inconsistent with the patterns of income and wealth we find. The data shows that there is substantial wealth mobility over individual's lifetimes that does not correspond to income mobility.

What is then responsible for differences in wealth accumulation across individuals? We simulate counterfactual wealth trajectories for each individual to gauge the importance of differences in saving rates and returns.[5] Our results show that portfolio-implied differences in returns are also insufficient to account for wealth dynamics, highlighting the importance of within-asset class heterogeneity in returns (Fagereng, Guiso, Malacrino, and Pistaferri 2020). Instead, we can rationalize these

[4] See, among others, Benhabib, Bisin, and Luo (2017); Stachurski and Toda (2019); De Nardi, Fella, and Paz-Pardo (2020); Sargent, Wang, and Yang (2021), for discussion of top wealth inequality.

[5] Because the data on asset transactions is only available starting 2006, it is not yet possible to conduct a full empirical decomposition of the contribution of income, savings rates and returns for the entire working life of a cohort, which is the focus of this paper.

differences through heterogeneous saving rates, but the magnitude of the differences is implausible. This requires 40 percent of the population to maintain a saving rate of over 40 percent, which is at odds with the evidence that the saving rate out of non-capital income is 7 percent across the wealth distribution in Norway (Fagereng, Holm, Moll, and Natvik 2019). Taken together, our results imply that persistent idiosyncratic differences in returns and saving rates are required to rationalize saving behavior, as no single element can generate the joint patterns of income and wealth we observe.

Finally, we contrast the role of individuals' circumstances, including parental wealth and education, in predicting full wealth rank histories.[6] We find an important and nonlinear role for family background. Individuals born to parents at the top of the wealth distribution are almost 30 percentage points more likely to be persistently at or near the top of their own generation's distribution, compared to those born to parents at the bottom of the distribution, who are more likely to be persistently poor throughout their lives. For individuals who rise or fall through the distribution, education is the main predictor of their evolution. Highly educated individuals are markedly more likely to rise through the wealth distribution as they age. By contrast, even after controlling for their parental background, those without post-secondary education are 5 to 10 percentage points more likely to be fallers than those with at least undergraduate degrees.[7]

These results provide a novel approach to studying intergenerational mobility in terms of entire life-cycle histories. We find declining intergenerational mobility along the life cycle, so that the wealth ranks of individuals move closer to their parents' ranks as they age. Not only does this mirror the trend in intragenerational mobility we document, but we find that the same individuals drive both population trends. As risers

[6]Our exercise moves beyond standard measures of intergenerational mobility that compare the rank of different generations at a similar point in their life cycle, thus, relying on a snapshot of their wealth trajectory to infer mobility (see, for example, Chetty et al. 2014; Fagereng, Mogstad, and Rønning 2021). We instead ask whether individual characteristics can predict complete life-cycle histories.

[7]Our results also complement those in Huggett, Ventura, and Yaron (2011), who study lifetime inequality using a model-driven approach. Although we focus on mobility, we both find important roles for human capital and initial conditions, including the initial wealth level of individuals.

rise and fallers fall, their reversals of fortune drive increasing intragenerational and decreasing intergenerational mobility.

The main methodological contribution of the paper is to propose a data-driven approach to summarizing heterogeneous mobility in large-scale datasets. The agglomerative hierarchical clustering algorithm we employ works by recursively grouping individuals with similar wealth-rank histories.[8] This process results in a global hierarchy of clusters that minimizes the distance between the paths taken by individuals in each group, making use of the whole vector of realized wealth ranks. Crucially, our methodology allows us to characterize mobility patterns without resorting to a single summary statistic; it also does not require us to specify which observable characteristics determine the groups or to rely on a specific parametric model for the evolution of wealth. To the best of our knowledge, this approach has not been applied to the study of mobility prior to this paper.

Our study of wealth mobility provides a variety of new measures that help us understand when and why different people save and accumulate wealth. These are useful inputs for models of saving over the life cycle, which are often not validated against the observed dynamics of wealth and, if they are, stop at fixed-periods transitions or focus exclusively on the very rich. We show that the joint dynamics of income and wealth are inconsistent with standard models of buffer-stock savings, and that no single additional factor (earnings risk, portfolio heterogeneity, savings rates) can, in isolation, account for the differences in wealth accumulation over the life-cycle. These results reinforce evidence on the crucial role of persistent return heterogeneity for wealth mobility, even for individuals away from the top of the wealth distribution.

[8] See Hastie, Tibshirani, and Friedman (2009, ch. 14) for an introduction to clustering; Borysov, Hannig, and Marron (2014), and Egashira, Yata, and Aoshima (2024) derive asymptotic properties of hierarchical clustering.

Related literature. We provide new evidence on wealth mobility along the life cycle, not only measuring the degree of persistence in individuals' positions in the wealth distribution but also the ways in which individuals move by characterizing their typical trajectories. In doing so, we complement an extensive literature on the dynamics of earnings over the life cycle (see, for instance, Arellano, Blundell, and Bonhomme 2017; De Nardi, Fella, and Paz-Pardo 2020; Guvenen, Karahan, Ozkan, and Song 2021; and Guvenen, Kaplan, Song, and Weidner 2022) and across generations (see, for instance, Solon 1992; Chetty, Hendren, Kline, Saez, and Turner 2014; Chetty, Grusky, Hell, Hendren, Manduca, and Narang 2017; and Halvorsen, Ozkan, and Salgado 2022). Relatedly, Hurst, Luoh, Stafford, and Gale (1998) study how saving behavior differs over a decade by race, education, household demographics, and initial wealth using the Panel Study of Income Dynamics. Instead, we study mobility with wealth trajectories over 25 years without conditioning on specific variables.

The dynamics of wealth mobility that we document provide much richer information than observing the cross-sectional wealth distribution at a moment in time. We argue that correctly replicating these dynamics is key to correctly capturing people's savings motives and, hence, to study a wide array of economic problems, from the pass-through of monetary policy (for example, Kaplan, Moll, and Violante 2018; Auclert 2019), to the crowding-out effects of social security (for example, Samwick 2003; Scholz, Seshadri, and Khitatrakun 2006; Blau 2016), and the optimal taxation of capital (Guvenen, Kambourov, Kuruscu, Ocampo, and Chen 2023; Phelan 2025).

We also contribute to the literature on intergenerational mobility in wealth (see, for instance, Charles and Hurst 2003; Boserup, Kopczuk, and Kreiner 2017; Adermon, Lindahl, and Waldenström 2018; Fagereng, Mogstad, and Rønning 2021; and Sabelhaus 2024). Our results show that single-age snapshots of intergenerational wealth persistence, a standard measure in the literature, imperfectly capture the link between individuals' wealth ranks and their parents' by not accounting for their full wealth

histories.[9] We show that both parental background and their own positions in the wealth distribution near the beginning of their work-life have long-lasting impacts on wealth trajectories. Intergenerational mobility declines over the life cycle as children's relative positions in their own generation converge toward those of their parents.

Our analysis is made possible by longitudinal data characterizing the distribution of wealth histories compiled by Statistics Norway. Observing individuals over long periods of time is crucial for studying the nature of wealth accumulation and prior contributions have used this data to investigate the role of return heterogeneity (Fagereng, Guiso, Malacrino, and Pistaferri 2020), differences in saving behaviors (Fagereng, Holm, Moll, and Natvik 2019), the importance of gifts and inheritances for lifetime resources (Black, Devereux, Landaud, and Salvanes 2022), and the relationship between wealth and lifetime income (Black, Devereux, Landaud, and Salvanes 2023). In related work, Ozkan, Hubmer, Salgado, and Halvorsen (2023) focus on the drivers of wealth accumulation among the wealthiest 0.1 percent at age 50 looking backward at their lifetime trajectories. We complement these papers by characterizing the life-cycle paths of individuals across the entire wealth distribution, including those with rising, falling, and stable paths. Our findings therefore contribute to our understanding of wealth inequality and mobility beyond the dynamics of wealth accumulation at the very top.

The clustering method we employ constitutes a feasible way to study trajectories of longitudinal outcomes, such as mobility, in large panel datasets. It also allows us to decompose commonly used summary measures of mobility, such as the OLS coefficient in a rank-rank regression. Similar approaches have been used in sociology to summarize mobility between discrete states (Dijkstra and Taris 1995; McVicar and Anyadike-Danes 2002; Dlouhy and Biemann 2015). In economics, clustering has been used to analyze sorting and transitions in the labor market (see, among others, Bonhomme, Lamadon,

[9] Closer to our focus in this paper, Shiro, Pulliam, Sabelhaus, and Smith (2022) consider wealth mobility during individuals' "prime wealth accumulation years," between their early thirties and late fifties. See Deutscher and Mazumder (2023) for a recent survey of different methodologies.

and Manresa 2019; Gregory, Menzio, and Wiczer 2021; Humphries 2022; and Ahn, Hobijn, and Şahin 2023) and to identify latent heterogeneity, as in Lewis, Melcangi, and Pilossoph (2021). Many of these applications use variants of K-means clustering, whose asymptotic properties are derived in Bonhomme and Manresa (2015) and Bonhomme, Lamadon, and Manresa (2022). Relative to these methods, our approach provides a global hierarchy of partitions that facilitates studying within cluster heterogeneity without imposing computational burdens in the analysis of large datasets. Although hierarchical clustering is our preferred approach, we show in Section 9 that our main results hold with K-means clustering.

2. Data: A panel of wealth histories for the Norwegian population

We employ data from the Norwegian tax registry between 1993 and 2017 and its associated population characteristics files. We are able to link these various datasets at the individual and household levels using unique (anonymized) identifiers. The resulting data contains information on wealth (net worth), assets, debt, income, and a variety of individual characteristics. We provide a detailed description of the data in Appendix A. We report monetary values in 2019 US dollars.

The coverage and properties of the Norwegian administrative data sets it apart from survey and administrative data available in other countries and makes it uniquely suited to the study of wealth mobility over the life cycle.[10] We start by highlighting the key strengths of our data.

First, Norway has recorded wealth in its tax returns since 1993, providing us with twenty-five years of observations. This long panel allows us to track individuals through important phases of their lives. Tracking individuals is crucial to understand

[10]The quality and detail of this data have proven useful in a variety of studies. More information on the Norwegian administrative wealth data can be found in Fagereng, Guiso, Malacrino, and Pistaferri (2020), Fagereng, Mogstad, and Rønning (2021), and Fagereng, Holm, and Natvik (2021). Additionally, Blundell, Graber, and Mogstad (2015) provides a detailed discussion of income tax records.

mobility over long horizons and to differentiate the life-cycle trajectories experienced by individuals, which we do using the clustering procedure described in Section 4.

Second, the Norwegian income and wealth tax records cover the entire population. We therefore construct accurate measures of an individual's rank in the wealth distribution, within cohorts and the population at large. Furthermore, the data covers individuals at the very bottom and top of the distribution, who are typically difficult to capture in survey data.[11] Moreover, most of the components of income and wealth are third-party reported and are not top- or bottom-coded, eliminating concerns about measurement error from self-reporting and censoring that are common in survey data.

Third, we are able to link individuals within households and across generations, as well as to their demographic and educational information. This wealth of information lets us link the trajectories of wealth mobility to the individual circumstances that help determine them, such as parental background and educational attainment.

2.1. Wealth and asset data

We observe each individual's assets, debt, and net worth, as reported in their wealth tax return from which we are able to construct the market value as described in Appendix A. These are individual returns, where the value of assets jointly owned by a couple is split equally between each partner. We focus our analysis on wealth at the individual level, but we also report robustness results for wealth at the household level. We also observe the value of various asset classes included in individuals' wealth tax returns. However, the tax returns do not include transactions within classes, preventing us from computing asset returns at the individual level.

[11]This problem has led to methods that oversample the tails of the distribution. These methods are ill-suited to the focus of our study. For example, the U.S. Panel Study of Income Dynamics oversamples lower income households (the Survey of Economic Opportunity households), while the Survey of Consumer Finances oversamples wealthier households. Researchers often resort to ad hoc methods to build more accurate measures of the upper tail of the wealth distribution, for example, by augmenting the Survey of Consumer Finances with the Forbes 400 list of the 400 richest Americans or estate tax data (see, for example, Vermeulen 2016). Davies and Shorrocks (2000) provide an extensive review of these methods.

The single largest asset for most individuals is housing. We adjust housing values using the adjustments reported in Fagereng, Holm, and Torstensen (2020) for owner-occupied housing, secondary housing, and cabins (holiday homes), treating condominiums separately from other properties.[12] The other asset classes included in the tax returns are vehicles, public and private equity, and safe assets. Vehicles includes cars and boats. Public equity is defined as directly owned stocks that are traded on the Norwegian Stock Exchange. Private equity includes the value of business assets and unlisted stocks. Our measure of safe assets includes government bonds, checking accounts, and shares in money market and mutual funds.[13] An individual's wealth tax return also lists foreign non-property assets and a residual class that includes hard-to-value assets—such as jewellery and paintings. We include these two classes in wealth, but do not report results for them in the paper. The asset classes we observe correspond to those referenced in Fagereng et al. (2020), from where we obtain average returns by asset class.

Two types of assets are missing from our data. First, assets individuals obscure from the tax authority. Although third-party reporting should minimize opportunities for tax evasion, these assets are not observed in tax data by definition.[14] Second, we lack information on pensions, including employer-provided pension plans. Private pensions represent less than 20 percent of all pensions in Norway, with pay-as-you-go public pension entitlements making up the remaining (Ozkan et al. 2023).[15]

[12]We thank Fagereng, Holm, and Torstensen for providing us with updated adjustment values covering our sample period.

[13]We view the last of these items as less safe than government bonds and deposits; however, data restrictions prevent us from considering an alternative definition where we pool this with public equity.

[14]We conjecture that tax evasion does not pose a large measurement issue for wealth ranks for two reasons. First, because the individuals hiding their assets abroad are likely to rank high in the observed wealth distribution, and second, most evasion is likely to be monotonic in post-evasion wealth rank.

[15]Pay-as-you-go pensions are annuities, and do not constitute wealth that can be accessed or pledged as collateral by working-age households. Hence, they are not included in wealth tax records, and we do not include them in our baseline measure of wealth. See Fagereng, Holm, Moll, and Natvik (2019, Appendix C.6) for details on the Norwegian Public Pension system and the imputation of Public Pension wealth.

2.2. Additional data: income and demographics

We also use high-quality information from individual income tax records, analogous to the wealth tax records described above. These records allow us to study gross and net income. Furthermore, we observe several components of income, including wage earnings, self-employment earnings, capital income, and transfers from social assistance programs. We relate income and wealth profiles in our analysis.

We also have access to detailed information on individual education levels and fields of study, according to the Norwegian Standard Classification of Education (NUS2000). This classification provides nine levels of education, ranging from no education to post-graduate PhD level, as well as 350 fields of study.

In addition, we merge several key demographic variables. These include individual attributes such as date, place, and sex at birth, as well as parents' identifiers, date of death, and immigration status. Finally, we observe the individuals' civil status, as well as their cohabitation status for each tax year as recorded by Statistics Norway (SSB).

2.3. Sample selection

We begin with the universe of Norwegian tax residents between 1993 and 2017. We then create a broad cross-cohort sample with individuals born after 1905 (Norwegian independence) and before 1990. We also exclude individuals with short panels, specifically, those who ever emigrated from Norway and those who either immigrated after the age of 25 or who arrived after 2011.

Our main sample is the 1960–64 birth cohort. We use this sample to calculate the

within-cohort wealth ranks that are our main object of interest.[16] This birth cohort is observed for a significant fraction of their work lives, starting in their early thirties. In addition, this cohort is not affected by the 1959 changes in the compulsory school age. This reform was not implemented uniformly across place and time; see Black, Devereux, and Salvanes (2005); and Bhuller, Mogstad, and Salvanes (2017).

Finally, we further restrict the sample to ensure that it is balanced over our 25-year panel when analyzing complete wealth trajectories, leaving a total of 279,002 individuals. This eliminates attrition driven by migration and mortality. Increasing mortality in late middle age drives a large share of this sample selection criterion.

3. Measuring wealth mobility along the life cycle

We now turn to using this panel of wealth histories to measure wealth mobility over the life cycle. We find substantial wealth mobility that increases as individuals age.

3.1. Wealth ranks and the Norwegian distribution of wealth

For our cohort of interest, born 1960–64, we construct yearly individual ranks of wealth (net worth) using the unbalanced subsample from 1993 to 2017. Formally, given individual i's wealth in year t, $w_{i,t}$, we compute ranks within the cohort for each year as

$$y_{i,t} = 100 \times F_w \left(w_{i,t} \mid t, i \in \text{Birth Cohort: 1960–64} \right), \tag{1}$$

[16] To illustrate the value of our data, our sample selection criteria yield 292,222 individuals in the 1960–64 birth cohort. By contrast, and before imposing any additional restrictions, there are only 1,463 unique households in the Panel Study of Income Dynamics (PSID) in the same birth cohort, and their wealth is only observed consistently since 1999, implying an average of just six consecutive observations (ten years) for any given head of household. Hurst, Luoh, Stafford, and Gale (1998) study a similar question to us using three observations from the PSID between 1984 and 1994, pooling all cohorts together to increase the sample size. They thereby capture both the effects of the age-profile and mobility among peers.

where F_w denotes the empirical cumulative distribution of wealth. Crucially, all comparisons use other individuals in the same cohort as the reference group. As a result, our rank measure is not affected by cross-cohort or cross-age comparisons.[17]

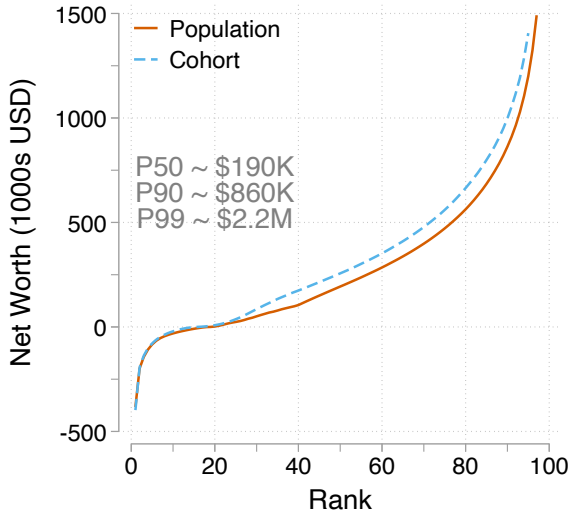
Ranks have several attractive properties over alternative transformations of net worth. They are well defined for individuals with negative or zero net worth and easily interpretable in terms of relative mobility in both levels and changes. In addition, they capture diminishing marginal gains from wealth for individuals because they compress the right tail of the distribution. For these reasons, we use ranks to study mobility—a common choice in the intergenerational mobility literature (see, for example, Chetty, Grusky, Hell, Hendren, Manduca, and Narang 2017).[18]

In Figure 1, we report wealth in US dollars by rank. Wealth in Norway is very unequally distributed. For reference, the 90$^{\text{th}}$ percentile of wealth in Norway is close to 860,000 US dollars, higher than in the US where it is 620,000 dollars (Smith, Zidar, and Zwick 2022). We observe that changes in ranks are associated with significant changes in wealth levels. For instance, moving from percentile 50 to 60 means going from 190,000 to 250,000 US dollars of net worth. The only part of the distribution in which rank changes do not translate into substantial movements in net worth is the narrow window around zero wealth (15$^{\text{th}}$–20$^{\text{th}}$ percentile). Moreover, changes in rank reflect meaningful differences in wealth even at younger ages, as we show in Appendix B.

[17] Importantly, doing this also purges ranks from time effects varying by age. For instance, all members of our sample experience the 2008 global recession at approximately the same age. We consider cross-cohort differences in life cycle wealth accumulation by analyzing the mobility patterns of the 1965–69 birth cohort in Section 9. See, for example, Gale, Gelfond, Fichtner, and Harris (2021); and Paz-Pardo (2024), for cross-cohort changes in wealth accumulation in the US.

[18] We also report results for absolute—rather than relative—wealth mobility from trajectories of (log) wealth and for alternative measures of the position of individuals in the wealth distribution, which deliver similar qualitative findings to our main analysis. These results are described in Section 9.

FIGURE 1. Norwegian wealth distribution in 2014, cohort ages 50–54



Notes: The figure shows the inverse CDF of the Norwegian wealth distribution in 2014 for the population at large (solid-orange line) and for the 1960–64 birth cohort (dashed-blue line) who are ages 50–54. Numbers are wealth holdings in 2019 US dollars by rank.

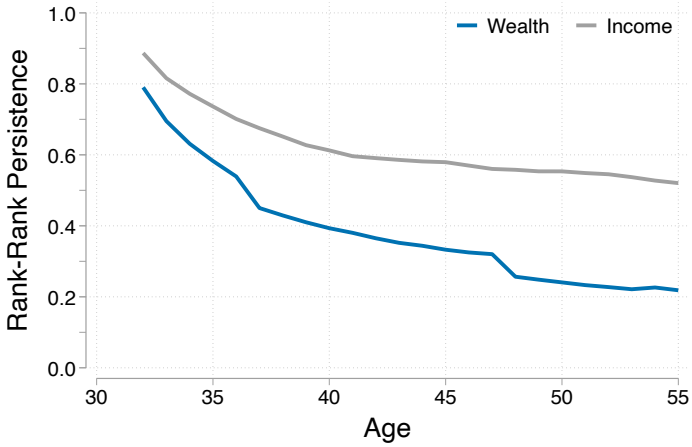
3.2. Rising wealth mobility along the life cycle

To what extent do individuals transition across the cohort wealth distribution over their lives? We measure *intragenerational* wealth mobility with the auto-correlation of wealth ranks as individuals age relative to their initial position in their cohort's wealth distribution. We compute these auto-correlations by regressing the rank in each year of our sample, $y_{i,t}$, on their rank when first observed in 1993, $y_{i,1993}$.[19] This provides age-varying measures for the *persistence* of individuals' position in the wealth distribution, with a lower persistence implying more mobility. As a point of comparison, we similarly measure the persistence of individuals' rank in their cohort's income distribution.

Figure 2 shows that intragenerational persistence in wealth decreases as individuals grow older, evidencing rising wealth mobility along the life cycle. Thus, individuals

[19] Formally, we estimate $y_{i,t} = \alpha_t + \rho_t y_{i,1993} + u_{i,t}$. Chetverikov and Wilhelm (2023) derive the asymptotic distribution of ρ_t which is non-standard. We repeat the analysis using the Shorrocks (1978) index of mobility as an alternative to rank auto-correlations in Appendix B and obtain very similar results.

FIGURE 2. Intragenerational persistence of wealth ranks



Notes: The figure plots the persistence measures of wealth and income against age, corresponding to the auto-correlation of wealth and income ranks, respectively, $y_{i,t}$, with their value in 1993, $y_{i,1993}$.

experience increasingly large (cumulative) changes in rank that persistently change their position in the wealth distribution.[20] Most of these changes take place early on, when persistence declines rapidly, and are less frequent from age 40 onward, when persistence stabilizes around 0.3 before decreasing again a decade later to 0.2.[21] By contrast, although the persistence of income ranks also decreases with age, income mobility is markedly more limited than wealth mobility. Income persistence stabilizes around 0.55 by 40, showing that individuals' position in the income distribution tends to be mostly stable throughout their work life, while their wealth position changes.

These results establish a trend of increasing wealth mobility, but they remain silent over how broad-based this mobility is. The persistence measure in Figure 2 collapses the myriad of wealth trajectories experienced by individuals into a single aggregate time series. Some of these trajectories correspond to individuals with relatively stable ranks

[20] In Section 8.2, we find the opposite trend for intergenerational wealth mobility: the relationship between individuals' wealth ranks and their parents' ranks in 1993 grows stronger with age.

[21] Our results are consistent with Boserup, Kopczuk, and Kreiner (2018) who find a rank-correlation between age 18 wealth and age 45 wealth of 0.22 in Danish administrative tax records. Using broadly similar definitions, Shiro, Pulliam, Sabelhaus, and Smith (2022) estimate greater persistence in wealth ranks (0.59) for the US over the same prime wealth accumulation years we study.

(and low mobility), and some to individuals who undergo large changes, rising or falling through the wealth distribution. Put another way, Figure 2 tells us that rank changes occur, but does not tell us the shapes of individuals' typical wealth histories, how stable or changing they are, or what the usual timing and magnitude of rank changes are.

4. Grouping life-cycle trajectories of mobility

To understand the underlying life-cycle patterns of wealth behind increasing mobility—who is moving? how much? and when?—as well as the economic mechanisms that shape them, we move to analyzing the distribution of wealth trajectories. We do this by clustering the 1960–64 birth cohort into groups with typical life-cycle trajectories that capture the variation in wealth histories, as we explain next. We show that these typical trajectories capture persistent differences in life experiences, describing a pattern of segmented wealth mobility. In a second step, we study the characteristics of each group and relate them to their wealth trajectories.

The trajectory of an individual through the wealth distribution over the duration of the panel is described by the vector of ranks

$$\mathbf{Y}_i = (y_{i,1993}, y_{i,1994}, \cdots, y_{i,2016}, y_{i,2017}) \in [0, 100]^{25} , \tag{2}$$

where $y_{i,t}$ is the within-cohort wealth rank of an individual as defined in equation (1).

The distribution of $\mathbf{Y}_i$ across the population is a high-dimensional object and therefore we proceed by reducing it to a small number of groups. We recover a set of $G > 1$ disjoint groups of individuals, so that each individual i is assigned to one of these groups, $g_i \in \{1, \ldots, G\}$. This induces a partition $\mathcal{G}_G = \{g_i\}_{i=1}^N$ over the set of individuals.

Specifically, we define groups of individuals with similar life cycles of wealth

mobility using an *agglomerative hierarchical clustering* algorithm.[22] Hierarchical clustering works recursively, starting from the lowest level of hierarchy, where $G = N$ and each observation is assigned to its own group, and sequentially combining (or agglomerating) one pair of groups in each iteration. This process results in a hierarchy of partitions ranging from $G = N$ to $G = 1$. At each level of hierarchy $G > 1$, the algorithm creates the partition at the next level $\mathcal{G}_{G-1}$ by combining the two groups with the lowest dissimilarity. We use Ward's method to agglomerate clusters and adopt the total within-cluster variance as the dissimilarity metric:[23]

$$\underset{g,g' \in G,\, g \neq g'}{\operatorname{argmin}} \quad d\left(g, g'\right) = \sqrt{\frac{2N_g N_{g'}}{N_g + N_{g'}}} \times \left\| \bar{\mathbf{Y}}^{g} - \bar{\mathbf{Y}}^{g'} \right\|_2 , \tag{3}$$

where g and g' are disjoint groups, N_g denotes the number of observations in group g, and $\bar{\mathbf{Y}}^{g}$ is the centroid (average) of the observations in group g.[24]

Crucially, we use the complete vector of ranks $\mathbf{Y}_i$ when grouping the life-cycle trajectories of mobility. Doing this has the key advantage that we do not need to assume that only a subset of the elements of $\mathbf{Y}_i$ are informative, as is the case when focusing on transitions over fixed horizons. Neither do we need to reduce the dimensionality of the object of interest to a single summary statistic, such as the rank persistence in Figure 2.

[22] It is possible to include other outcomes or covariates that differ across groups, such as their income or portfolio composition. However, in practice, doing this introduces more noise than additional information; accordingly, we focus on wealth ranks in our main analysis.

[23] Alternative specifications of the dissimilarity metric, including maximum or median distance, are also possible. See Humphries (2022) for another application of Ward's method in the context of Sequence Analysis, where it is used to cluster panel data with discrete states. We produce our Agglomerative Hierarchical Cluster Tree using Matlab, see https://www.mathworks.com/help/stats/linkage.html.

[24] Given a set of underlying groups $G^\star$, a classifier is asymptotically consistent if, as the length T of observed trajectories increases, the classifier does not produce mixtures over these groups until it is asked to provide a partition into $G < G^\star$ groups. Borysov, Hannig, and Marron (2014) show this is the case for Ward's method as either $T/N \to \infty$, with T growing faster than N, or only $T \to \infty$, when the true group specific densities are jointly normal. For fixed population size N, Egashira, Yata, and Aoshima (2024) strengthen these results for arbitrary densities. These results highlight the importance of long panels such as ours that provide sufficiently long enough trajectories to distinguish among groups.

Selecting the number of groups. A key feature of this approach, which distinguishes it from other commonly used algorithms like K-means, is that we do not need to pre-specify the number of groups to study. Instead, implementing the algorithm recovers a complete hierarchy of nested groups. This makes it straightforward to study typical trajectories for any number of groups and to decompose the heterogeneity within each group by exploiting the nested structure. Therefore, we select the number of groups for our main analysis after we obtain the full hierarchy. We report additional exercises increasing the number of groups in Section 9, where we dissect the heterogeneity within each of the groups of our main analysis.

Selecting the number of groups used in our main analysis requires that we trade off two objectives: (*i*) having enough groups to represent the distribution of wealth rank histories, and (*ii*) having a parsimonious description of trajectories. Fewer groups provide a more easily interpretable picture of wealth mobility but may obscure relevant variation in the trajectories of the group's members. We find that setting $G = 4$ captures just over 50 percent of the variation in wealth ranks trajectories and use this as our baseline. See Appendix C.1 for details.

5. Segmented wealth mobility

Using agglomerative hierarchical clustering, we group individuals who experience similar trajectories through the wealth distribution as they age. The resulting typical wealth trajectories experienced by individuals in our four baseline groups capture very different mobility patterns. We show that these patterns are the results of segmented wealth mobility among the individuals in our sample: only some groups drive mobility, but they only move within the middle segment of the distribution.

5.1. Typical wealth trajectories over the life cycle

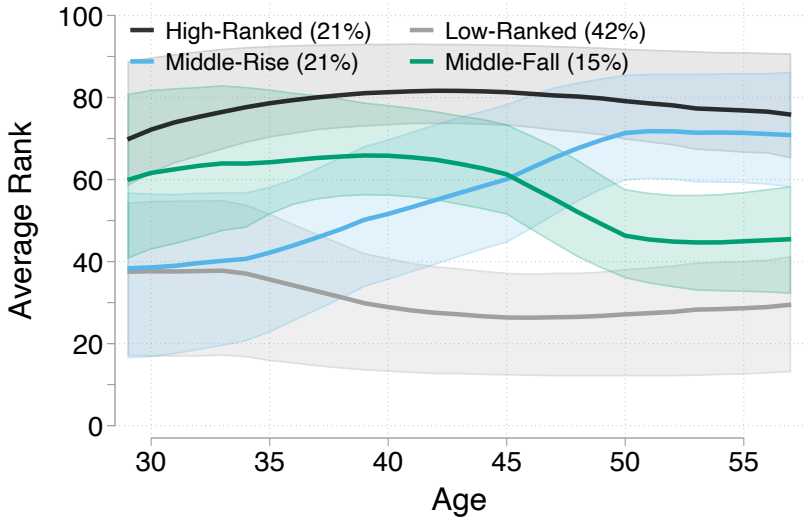
We begin by reporting the typical wealth rank trajectories of individuals in each of our four main groups in Figure 3. The typical trajectories have groups remaining at the bottom, in the middle, and at the top of the distribution throughout their life cycle, with the groups in the middle exhibiting rising and falling trajectories, respectively. Moreover, despite within-group heterogeneity, the interquartile range of the rank distribution for each group reveals that individuals' movements lie within segments of the wealth distribution. We interpret these patterns as evidence of *segmented wealth mobility*.[25]

Two groups of individuals, which we label "high-ranked" and "low-ranked," start their lives at the top or bottom of the wealth distribution and tend to stay there. They make up 21 and 42 percent of the cohort, respectively. This does not imply that their wealth rank is fully stable (as we show in Section 9.2) but that it tends to stay within the upper or lower segments of the wealth distribution, as made clear by the small changes in the interquartile range of the distribution of ranks of these groups.

The other two groups, which we label "middle-rise" and "middle-fall", correspond to the remaining 21 and 15 percent of the cohort, respectively. They stay in the central part of the distribution, but have, respectively, increasing and decreasing wealth rank trajectories. These trajectories lead them to overlap with the high-ranked and low-ranked groups by age 55. Crucially, the reversal of fortune experienced by these groups is the key driver behind the population trend of intragenerational mobility documented in Section 3. As we expand on later, the trajectory of the risers is the main driver of increased mobility along the life cycle and contrasts with the relatively lower mobility of the high- and low-ranked groups. In this way, mobility is not universal; rather, it is

[25]The same patterns arise under alternative clustering algorithms and for alternative outcomes, in particular, when clustering on variables that put more weight on differences in wealth levels, relative to the rank trajectories used in our main exercise. We discuss these alternatives in Section 9.

FIGURE 3. Life-cycle dynamics of wealth mobility



Notes: The figure plots the average wealth rank in each clustered group against individuals' age. The shaded areas correspond to the interquartile range of the rank distribution of each group for each age. All individuals belong to the 1960–64 birth cohort. The clusters are constructed using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3).

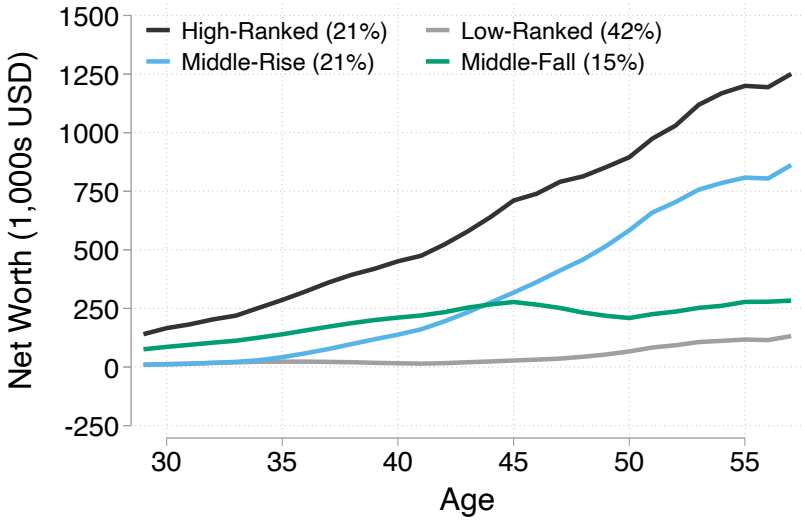
limited to segments of the population in the middle of the distribution.[26]

These typical wealth rank trajectories capture economically meaningful differences in the wealth trajectories of individuals, as we show in Figure 4. In particular, the large reversals in fortune experienced by risers and fallers reflect different trajectories of wealth accumulation and not spurious mobility generated by a compressed wealth distribution. We capture these differences in wealth trajectories because our long panel allows us to identify slow-moving patterns typical of wealth accumulation and because rank-differences in Norway reflect significant differences in wealth.

Figure 4 also gives several insights into the wealth mobility patterns of each group. For the high-ranked group, maintaining their position in the wealth distribution entails accumulating wealth as they age. Interestingly, the average rank among risers improves

[26]There are also other mobility patterns that are specific to smaller groups of the population. For instance, Ozkan, Hubmer, Salgado, and Halvorsen (2023) show that one-fourth of the wealthiest 0.1 percent at age 50 come from the bottom half of the distribution.

FIGURE 4. Average wealth by group



Notes: The figure plots the average wealth level in thousands of 2019 US dollars in each clustered group against individuals' age. All individuals belong to the 1960–64 birth cohort. The clusters are constructed using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3).

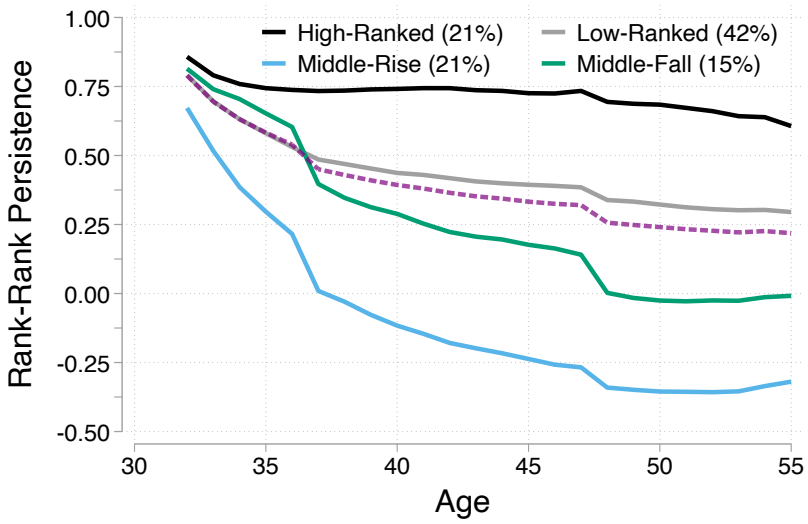
by accumulating net worth at a similar pace, but they start with zero net worth, on average, at age 30. By contrast, fallers have larger net worth than risers at age 30, but they accumulate wealth slowly, leading them to fall down the distribution. The effects of the Great Recession are clearly visible for this group, with a drop in their net worth around ages 45–50, when the recession hits the 1960–64 birth cohort.[27] Finally, the individuals in the low-ranked group have close to zero net worth, on average, for most of their work life and a very gradual increase starting at age 45.

5.2. Decomposing mobility patterns

The pattern of segmented wealth mobility in Figure 3 captures permanent differences in mobility between groups that, in turn, explain what (and who) is driving the trend of increasing intragenerational mobility described in Section 3. To see this, we decompose

[27] In Section 9.3 we show that the mobility trends we document are also present in other cohorts which were affected by the Great Recession and other macroeconomic events at different ages.

FIGURE 5. Intragenerational persistence across groups



Notes: The figure plots the intragenerational rank-rank persistence measure for our four main groups. The pooled cohort-level persistence measure is shown in dashed lines. The rank-rank persistence measure corresponds to the auto-correlation of wealth ranks, $y_{i,t}$, with their value in 1993, $y_{i,1993}$.

the persistence of wealth ranks by computing the rank-rank persistence measure separately for each of our main four groups.[28] We present the results in Figure 5.

Each group exhibits distinct patterns of wealth mobility. The high-ranked and low-ranked groups display higher levels of intragenerational persistence (lower mobility) than the population average. The persistence of individual ranks is particularly high and stable for the high-ranked group, barely decreasing until the group's members are more than 45 years of age. This is consistent with their observed wealth trajectories. The persistence of the relative wealth position of individuals in the low-ranked group is somewhat lower. Although these individuals tend to stay at the bottom of the wealth distribution, they experience more frequent movement within that (larger) segment of the distribution, in part reflecting the larger size of this group.

By contrast, mobility is higher for the groups of risers and fallers, reflecting large

[28] Formally, we compute the auto-correlation for each group g, ρ_t^g, estimating $y_{i,t} = \alpha_t + \rho_t^{g(i)} y_{i,1993} + u_{i,t}$.

changes in ranks over the life cycle that result in low and even negative rank autocorrelation. This pattern is strongest for the group of risers, for whom reversals of fortune happen relatively quickly (within the first 15 years of our panel) and whose rank-rank persistence measure eventually becomes negative.

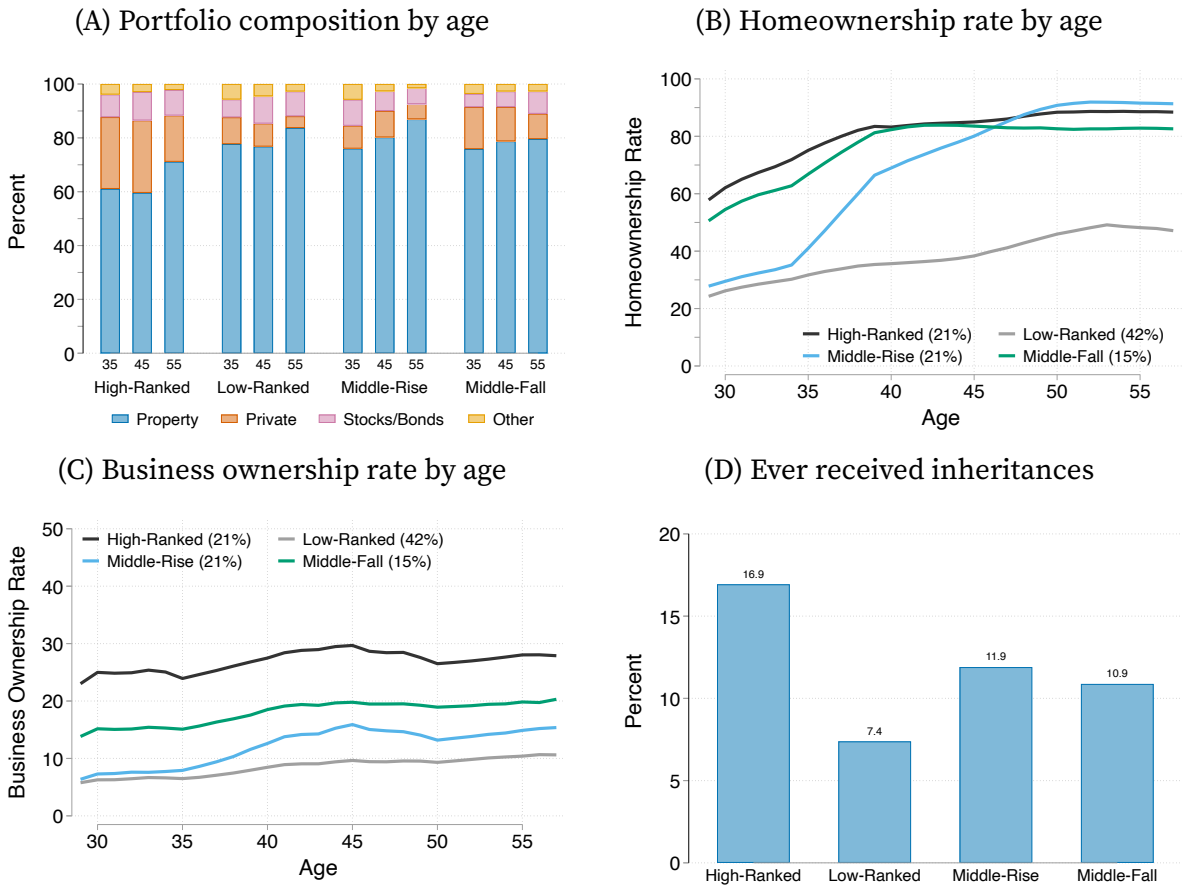
Overall, we find that the trend of increasing intragenerational mobility that we document in Figure 2 is not driven by broad-based mobility across the population or by large reshuffling across the distribution. Instead, the trend comes from a combination of stable groups at the top and bottom of the distribution and two groups undergoing relatively large transitions that are nevertheless contained to the middle segment of the wealth distribution. In essence, mobility is segmented, taking place for only some groups of individuals and within a section of the distribution.

6. Heterogeneity across groups: Portfolios, income, and households

We now turn to exploring the ex post life-cycle characteristics of each of the four main wealth mobility groups. We leverage the information available in the Norwegian Registry data to consider the main drivers of wealth accumulation over an individual's work life. Specifically, we look at portfolio composition, inheritances, sources of income, entrepreneurship, and marriage and divorce. These factors have been shown to be key determinants of wealth accumulation and wealth inequality (De Nardi and Fella 2017; Kuhn, Schularick, and Steins 2020; Hubmer, Krusell, and Smith 2021).

We quantify important differences between groups that help explain their divergent wealth accumulation and mobility patterns. However, the typical trajectories we recover encapsulate a wide variety of heterogeneous life-cycle events and choices. Consequently, we find no single factor that, in isolation, can account for these differences. We summarize the resulting patterns next and discuss the role of income, returns, and savings in Section 7.

FIGURE 6. Portfolio statistics by group



Notes: The figures present characteristics of the four main groups presented in Figure 3. Panel A reports the share of assets accounted for by property, privately held assets, financial assets, and other assets, defined as the total value of each asset class divided by the total assets within a group. Panels B and C plot, respectively, the share of individuals who are homeowners and who own business assets. Panel D reports the share of individuals in each group that have received inheritances by the age of 55.

6.1. Portfolio composition

Figure 6 describes the portfolio composition of each group. Panel A reports the share of assets accounted for by property, private business assets, financial assets (stocks, bonds, and bank accounts), and a residual category including vehicles and foreign assets.[29]

Property represents the majority of household portfolios across all groups; its share

[29]We find that the differences in net worth shown in Figure 4 mainly reflect differences in the accumulation of assets and not debt. The debt trajectories show much lower dispersion and levels than the asset trajectories, as we show in Figure C.4 in Appendix C.3.

increases slightly as homeownership rates increase. The increase in homeownership occurs in all groups, but the age profiles differ markedly (Figure 6B). Fallers are much more likely to be homeowners than risers up to age 45, when their homeownership rates converge. This finding is consistent with the higher initial asset position of fallers and the accumulation of wealth by risers out of labor income that we discuss below. As for the low-ranked group, their homeownership rate starts from 25 percent (a similar level as the risers), but it ultimately stalls at around 50 percent—well below the level of the other groups that have rates over 80 percent by age 55.

The portfolio composition of the high-ranked group stands out because of the lower share of property assets in their portfolio and correspondingly higher share of private business wealth, relative to the other groups. This pattern is driven by differences in the extensive margin of business operation (Figure 6C), as well as the group's higher shares of capital income and self-employment income, which we discuss below. Notably, the business ownership rate of the high-ranked group is close to 10 percentage points higher than that of the fallers and much higher than that of the risers; the risers' rate is never above 15 percent, although it increases as they age.[30]

Taken together, these patterns suggest that the ownership and operation of profitable business assets are crucial characteristics for many individuals who start and remain near the top of the wealth distribution; however, they do not play an equally important role for individuals rising through the middle of the wealth distribution. The accumulation of property is, instead, the main driver for the latter group, in line with the findings of Kuhn, Schularick, and Steins (2020) for the US.

Finally, wealth accumulation is also partially affected by intergenerational transfers, such as inheritances. Although we do not have direct information on all such transfers, the registry data does report whether individuals have received gifts or inheritances from a single source exceeding 470,000 NOK (about 45,000 US dollars) over their lives.

[30]We further describe heterogeneity in non-housing wealth accumulation across groups in Figure C.5.

Figure 6D shows that the share of individuals that received an inheritance by age 55 ranges from 7 in the low-ranked group to 17 percent in the high-ranked group. In the same context, Black, Devereux, Landaud, and Salvanes (2022) find gifts and inheritances represent a small fraction of an individual's net wealth at any point in time.[31]

6.2. Income trajectories

The income profiles of our main four groups are broadly consistent with the patterns of wealth mobility described previously, as we show in Figure 7. Nevertheless, the differences in income are smaller than the differences in wealth documented in Figures 3 and 4. Looking at income ranks, there is a 20-rank gap between the average profile of individuals in the high- and low-ranked groups. This gap is significant, but smaller than the 30- to 55-rank gap between the groups' wealth rank profiles.[32]

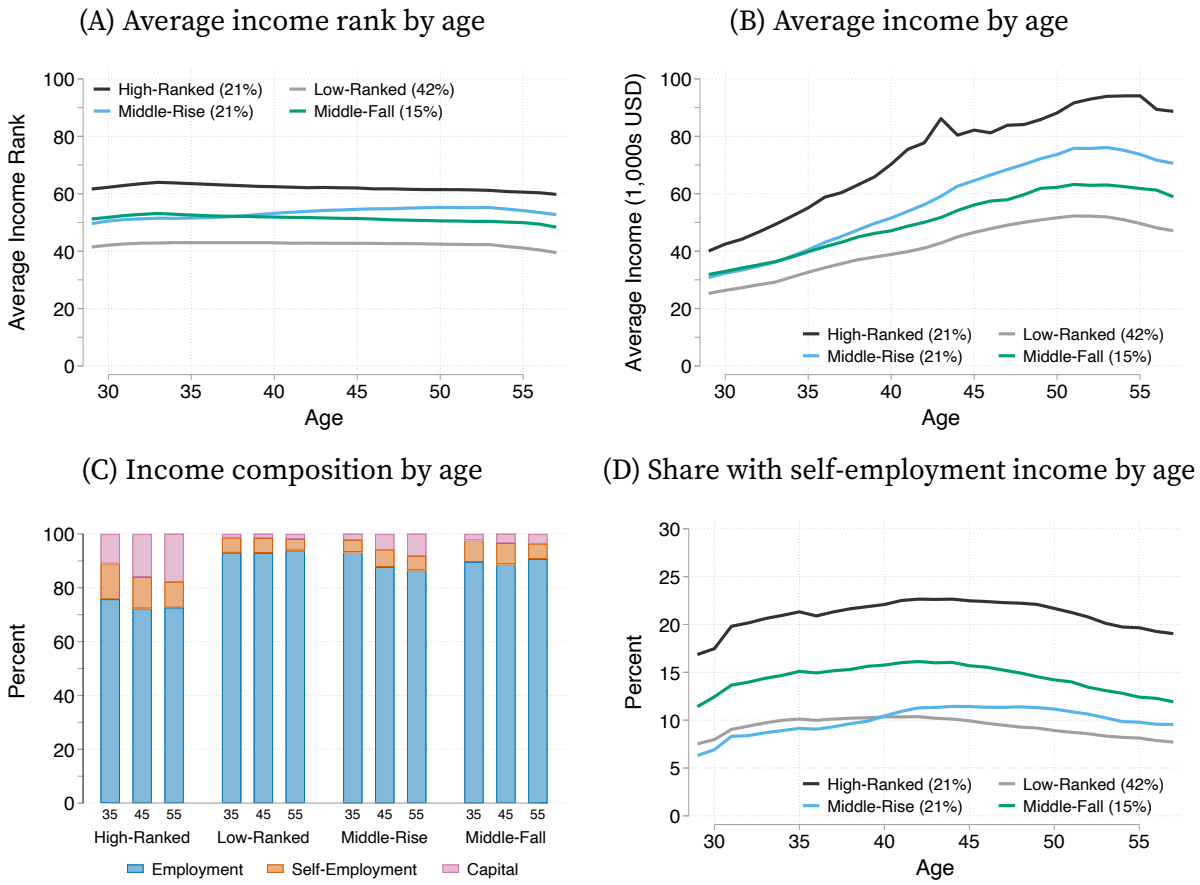
Despite smaller differences in income than in wealth, it is clear that the trajectory of income throughout individuals' working lives plays a relevant role in wealth mobility. Fallers begin with more wealth than risers and a similar level of income; yet, their incomes diverge after age 40. This widening income gap is coupled with a rapid increase in wealth for the risers and a reversal of their relative position with respect to the fallers. We show in Section 8 that these life-cycle differences in income are partly explained by higher educational attainment of risers (relative to all other groups), pointing to an important role for human capital accumulation in explaining wealth mobility.

In fact, taking into account *human wealth*, that is, the discounted value of future realized labor market income, closes the gap in initial financial wealth between risers and fallers. That is, the risers' higher income trajectories fully compensate for their

[31] Our results, and Black et al.'s, differ from Adermon, Lindahl, and Waldenström (2018) who emphasize the importance of intergenerational transfers in the form of bequests and gifts in Swedish data. But, consistent with the results we document below, they also find an important role for human capital.

[32] The income of those in the high-ranked group is, on average, 60 percent larger than that of the low-ranked group at age 30 and over 90 percent higher by age 55. The wealth of the high-ranked group is always at least ten times the wealth of the low-ranked group.

FIGURE 7. Income by group



Notes: Panel A plots the average income rank trajectories for the individuals in each of the four main groups presented in Figure 3. Panel B plots the average income in 2019 US dollars for each group. Panel C plots the share of each group's income accounted for by employee, self-employment, and capital income. Panel D plots the share of individuals in each group with self-employment income.

lower initial financial wealth. The sum of financial and human wealth remains relatively constant for the two groups until age 45 when their trajectories diverge, coinciding with the crossing pattern in financial wealth and the increase of non-property assets in the risers' portfolio documented above. We present the results in Figure C.9 and discuss them further in the Appendix.

Turning to the sources of income, Figure 7C shows that the vast majority of income

comes from wages and salaries in all groups.[33] The high-ranked group is the only one for which capital and self-employment represent a sizable source of income (up to 30 percent on average). Indeed, over 20 percent of the individuals in the high-ranked group have income from self-employment at age 45 (Figure 7D), more than double that of the low-ranked group. Similarly, the fallers are more likely to be self-employed than the risers (16 percent compared to 11 percent).

6.3. Household characteristics

We also track the civil status of individuals in each group. Figure 8A reports the share of individuals who are married, cohabiting, and single at several points during the panel. We find limited differences between groups. By age 45, over 80 percent of individuals in all groups are married or cohabiting, although individuals in the low-ranked group are slightly more likely to be single

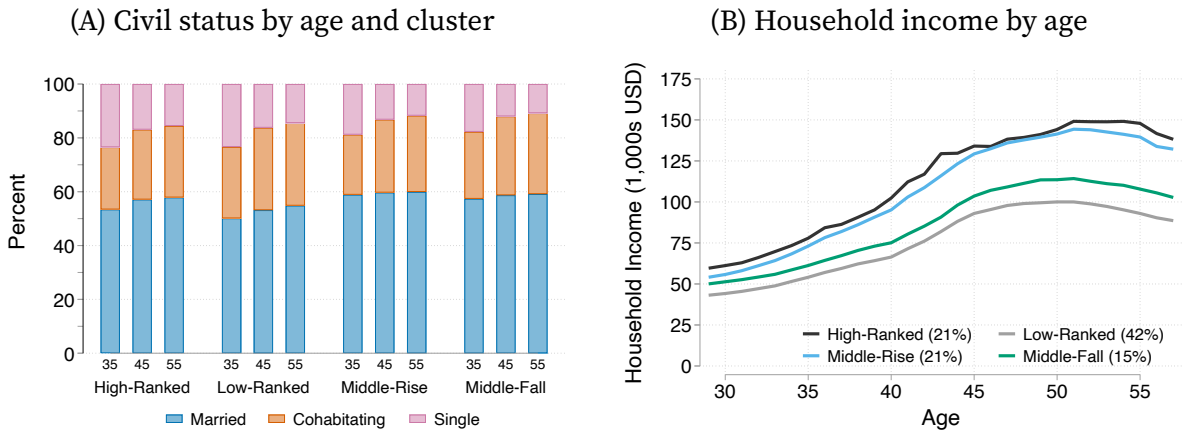
The high rates of marriage and cohabitation make the trajectories of household income that we report in Figure 8B relevant for the accumulation of wealth. To build household income, we construct households based on marriage and cohabitation using the complete population files; we assign to each individual in the 1960–64 birth cohort their household's income; and we equivalize using Organization for Economic Co-operation and Development's (OECD) equivalence scale based on the number of adults and children in the household.[34]

We find that the life-cycle profiles of household income reinforce, rather than reduce, the differences in individual labor income across groups (Figure 7B). We find a large gap in household income between the high-ranked and rising groups and the low-ranked

[33]We focus on sources of earned income, excluding income from public sector benefit programs. We show in Figure C.6 of Appendix C.3 that low-earning, low-wealth individuals are more likely to be beneficiaries of these programs. For example, 10 percent of the individuals in the low-ranked group receive unemployment benefits at age 40, and more than 20 percent receive disability benefits, compared to less than 5 and 10 percent, respectively, in the high-ranked group.

[34]We also compute household wealth ranks, reported in Appendix C.3, and show that our main results are robust to using household wealth instead of individual wealth in Section 9.

FIGURE 8. Household characteristics by group



Notes: The figures present characteristics of the four main groups in Figure 3. Panel A reports the share of married, cohabitating, and single individuals at ages 35, 45, and 55. Panel B plots average household income trajectories in 2019 US dollars.

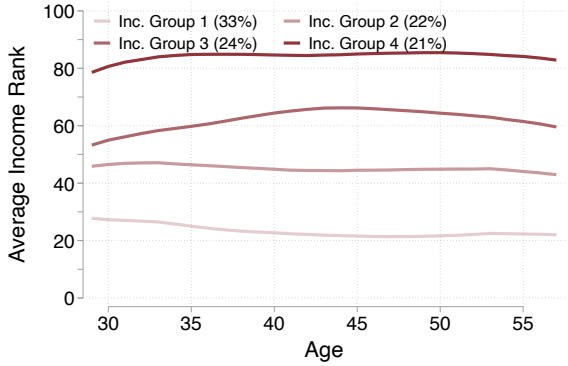
and falling groups. The patterns we find are consistent with higher-earning individuals marrying or cohabiting with higher-earning spouses, a pattern that is strengthened by households also sorting on the basis of their initial wealth and returns as documented by Fagereng, Guiso, and Pistaferri (2022). This assortative matching mechanism is particularly relevant for the individuals in the rising group, for whom household income is almost as high as that of the households of the high-ranked individuals.

7. The role of income, savings, and returns for wealth trajectories

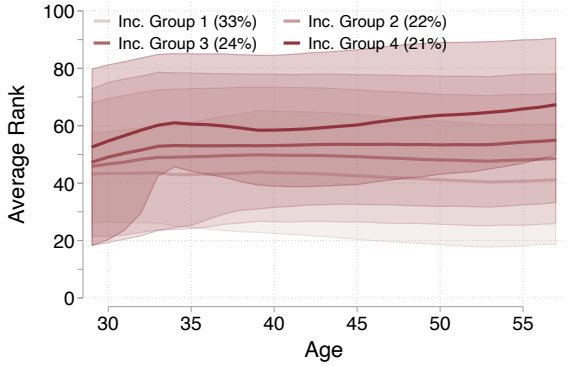
Despite revealing important differences in mobility that are correlated with individual factors, our results thus far do not disentangle how each of these individual factors contributes to the wealth mobility patterns of our four main groups. For example, differences in labor income could explain a substantial part of the wealth dynamics we observe as they cumulate over time. Alternatively, heterogeneity in saving rates or returns might be needed. In this section, we show empirically that the large divergence in wealth patterns requires differences in all three components across groups.

FIGURE 9. Grouping on income trajectories

(A) Life-cycle dynamics of income mobility



(B) Wealth mobility by income group



Notes: The figures present characteristics of four main groups recovered by clustering on income (analogously to results in Figure 3). All individuals belong to the 1960–64 birth cohort. The clusters are constructed using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3). Panel A plots the average income rank in each clustered group against individuals' age. Panel B plots the average wealth rank in each clustered group against individuals' age. The shaded areas in Panel B correspond to the interquartile range of the rank distribution of each group for each age.

7.1. Income dynamics, buffer stock savings and wealth mobility

We begin by determining whether the observed differences in income trajectories in our data are sufficient to explain the wealth patterns we have documented so far. To this end, we consider a common class of life-cycle buffer-stock savings models under homotheticity. These models are a useful benchmark for two reasons: first, they are extensively used in applications studying joint income and consumption dynamics (e.g., Zeldes 1989, Deaton 1991, Carroll 1992, and the large literature that followed them); second, they imply a tight and testable link between income and wealth mobility that will help us highlight the role of deviations from this benchmark in the form of heterogeneity in saving rates and returns.

In this class of models, finitely-lived agents choose how much to consume and save each period, earn a constant (and homogeneous) return on savings, r, and receive exogenous labor income that depends on a permanent component I (e.g., skill

heterogeneity) and idiosyncratic shocks z. We assume that income is log-separable in its permanent and transitory components, so that total income is $I \cdot y_t(z)$, where $y_t(\cdot)$ captures how efficiency units of labor reflect the value of z at age t.[35]

This setup implies that income profiles are, by construction, ordered by permanent income (skill). Thus, there is no mobility in income between groups, as the relative positions of individuals reflect differences in their permanent income components. This is consistent with our data, as we verify by clustering individuals based on their income trajectories. Figure 9 shows the results of this exercise: we obtain four lines in income ranks with no mobility between groups, pointing to the importance of permanent skill heterogeneity for income dynamics.

What are the implications of the buffer-stock savings model for wealth mobility? Under homotheticity in preferences, the properties of income are inherited by wealth and consumption, in a way that hierarchical clustering can recover. So, in the absence of other sources of heterogeneity, we should recover a pattern of no wealth mobility, with wealth groups corresponding to income groups. We make this precise in the following proposition and present the proof in Appendix D:[36]

PROPOSITION 1 (**Income and wealth trajectories in the buffer-stock model**). *Suppose agents with permanent income component, I, choose policy functions $c_t(a, I, z)$ and $a_{t+1}(a, I, z)$ to maximize utility from consumption, $u(\cdot)$, and bequests, $\upsilon(\cdot)$, in the following program*

$$\max_{\{c_t, a_{t+1}\}} \mathbb{E}_0 \left[\sum_{t=0}^{T} \beta^t S_t \left(u(c_t) + \frac{S_{t+1}}{S_t} \upsilon(a_{t+1}) \right) \right] \tag{4}$$
$$\text{s.t.} \quad c_t + a_{t+1} = (1+r)\, a_t + I\, y_t(z_t); \quad a_{t+1} \geq 0; \quad a_0 = 0;$$

where z is the agent's idiosyncratic income shock that follows a Markov chain, $z_t \in \mathcal{Z}$, with

[35] This assumption agrees with a wide range of approaches to modeling permanent income heterogeneity separating permanent traits (like education) from the dynamics of income (see, among others, Storesletten, Telmer, and Yaron 2004; Low and Pistaferri 2015; Guvenen, Karahan, Ozkan, and Song 2021; Guvenen, Kambourov, Kuruscu, Ocampo, and Chen 2023)

[36] The proof relies on Straub (2019), who proves that both policy functions and distributions over state variables are homothetic. Straub (2019) also shows that this extends to a setting with endogenous factor prices and the receipt of bequests in an overlapping generations economy.

age-dependent transition probabilities $\Pi_{zz'}(t)$ from state z to state z', and S_t is the probability of survival until age t.

If $u(\cdot)$ and $v(\cdot)$ are homothetic with the same constant elasticity, then policy functions are also homothetic, so that

$$c_t(a, I, z) = I \times c_t\left(\frac{a}{I}, 1, z\right) \quad \text{and} \quad a_{t+1}(a, I, z) = I \times a_{t+1}\left(\frac{a}{I}, 1, z\right) \quad \forall a \text{ and } \forall z \in \mathcal{Z}, \quad (5)$$

and therefore the wealth and consumption profiles are ordered in logs and ranks by permanent income.

Furthermore, given heterogeneity in permanent income, agglomerative hierarchical clustering is an asymptotically consistent classifier for the following latent groups

i. Groups partitioned by permanent income when clustering on income data;

ii. Groups partitioned by permanent income when clustering on wealth data;

iii. Groups partitioned by permanent income when clustering on consumption data.

Thus, the classification on income, wealth, and consumption profiles produces the same groups irrespective of the data on which clustering is based.

Before turning to the empirical implications of this result, we make two brief remarks.

REMARKS. *First, we allow for arbitrary concavity in the consumption function and do not require linearity of household decisions. Second, while this result is true in this broad, but stylized, class of models, it extends to richer frameworks as long as cross-group differences are meaningful, in the sense that differences in average outcomes between groups are larger than the variance of outcomes within groups.*[37] *We formalize this argument as a restriction on permanent heterogeneity in Appendix D.*

Proposition 1 effectively ties income and wealth mobility together. We contrast these implications with the data by comparing Figures 3 and 9. First, the fact that we obtain four parallel lines in Figure 9 confirms that income is more persistent than wealth even when we consider income trajectories in a flexible manner. Second, these income groups

[37] For example, when classifying on income with group-specific income risk, it must be that this risk is smaller than average differences in permanent income. See Meghir and Pistaferri (2011) for examples of this case, as well as Gourinchas and Parker (2002) and Meghir and Pistaferri (2004), who find that variability in residualized income is smaller than Mincer differences by education.

do not display significantly different wealth dynamics, suggesting that the tight link between income and wealth implied by homothetic buffer-stock models does not hold. Third, clustering on income and wealth recovers different groups of the population. We find the correlation between our wealth groups (Figure 3) and the income groups (Figure 9) to be relatively low, even though income trajectories differ between wealth groups (as shown in Section 6.2).[38] We conclude that the mobility patterns of standard buffer-stock saving models are inconsistent with the wealth mobility patterns we document and with the compression of wealth mobility for the income groups we recover.[39]

We view this as additional evidence in favor of richer models of heterogeneity, including heterogeneous savings behavior or heterogeneous returns that can generate meaningful non-homotheticities. Models that are approximately homothetic cannot rationalize the joint patterns of wealth and income mobility that our clustering method reveals. This finding is consistent with a large literature, based on life-cycle models, that shows that labor market income risk, even if it is modeled in a very flexible way, is not sufficient to generate the levels of top wealth inequality we observe in the data (Benhabib, Bisin, and Luo 2017; Stachurski and Toda 2019; De Nardi, Fella, and Paz-Pardo 2020; Sargent, Wang, and Yang 2021). Our results add to this by showing empirically that labor income risk is also insufficient to generate the patterns of wealth accumulation and mobility throughout the distribution.

7.2. Returns and savings rates

The previous results demonstrate that differences in income are not sufficient to explain the dynamics of wealth mobility that we document in the data. In light of this,

[38] In fact, the members of the wealth groups are well represented in each income group. For example, 18 percent of the high-ranked are in the lowest income group. This implies that individuals with very different income trajectories are present in each wealth group, as hinted in Figure 7A. We report the distribution of individuals across groups in Table C.1 in Appendix C.3.

[39] Further, the fact that risers and fallers have the same level of total (financial plus human) wealth until age 45 poses an additional challenge for explanations based solely on non-homotheticity in preferences.

we implement a set of counterfactuals to determine to what extent we could generate the observed trajectories of wealth accumulation if we allow for differences in income, returns, and savings rates between groups. In these counterfactuals, we assign to each individual their observed non-capital income and initial wealth, and construct counterfactual wealth trajectories under different return and savings rates that reflect their individual choices and the behavior of Norwegian savers.[40]

We leverage the variation in portfolio choices across individuals to compute counterfactual returns that reflect each individual's portfolio composition throughout their life (Section 6.1). We assign to each asset class its average return as reported in Fagereng et al. (2020, Table 3) and compute yearly portfolio-weighted returns for each individual of our sample. Figure 10 reports the average return profiles of our four wealth mobility groups. While differences in portfolio allocations generate differences in these implied returns, a striking fact is that for much of their life-cycle they imply that risers earn lower returns than fallers. We stress that these implied return differences are small relative to the dispersion in returns observed in the data (Fagereng et al. 2020).[41]

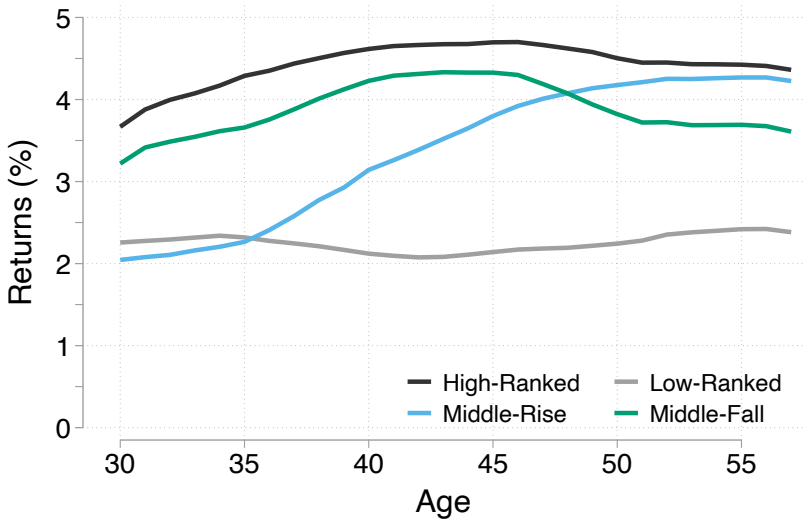
We compute counterfactual wealth trajectories under two *average* savings rates out of non-capital income net of taxes and transfers. The first, 7 percent, is obtained from Fagereng, Holm, Moll, and Natvik (2019) and represents *average active savings* in the Norwegian population.[42] The second, 45 percent, is an upper bound that serves as a reference to rationalize the behavior of the high-ranked and middle-rise groups. This number corresponds to the *marginal propensity to save* out of unearned income

[40] We also compute counterfactual wealth trajectories at the household level following the same procedure. They show the same patterns as for individual wealth.

[41] We observe assets separately from outstanding debt and, thus, account for leveraged returns in constructing the portfolio implied returns. Recall that we do not observe transactions within classes (Section 2), preventing us from incorporating heterogeneity in returns within asset classes.

[42] Fagereng et al. (2019) distinguish active saving rates out of income net of capital gains from saving rates including capital gains. Active saving rates are flat throughout the wealth distribution, while including capital gains implies increasing saving rates. These results for Norway are consistent with the results of Dynan, Skinner, and Zeldes (2004) for the United States.

FIGURE 10. Portfolio implied returns

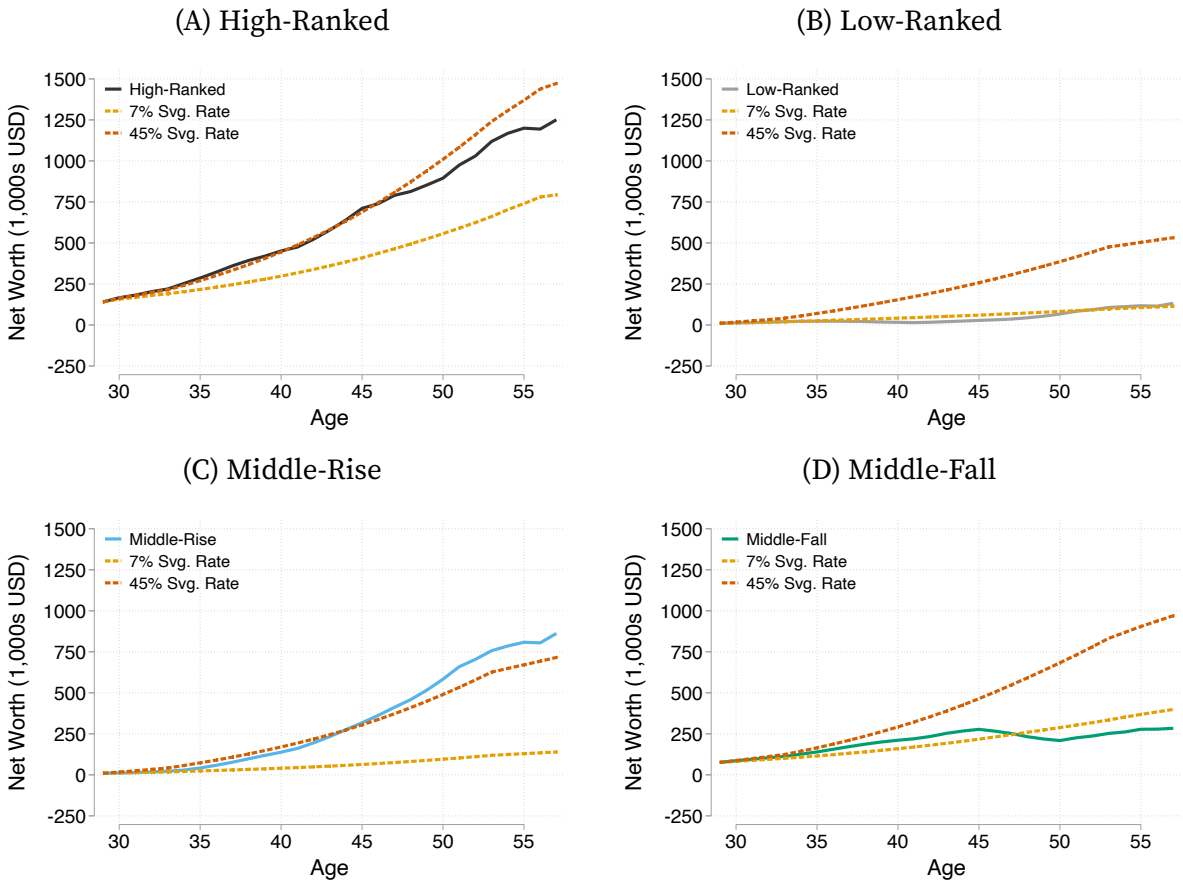


Notes: The figure presents average returns for the groups presented in Figure 3. We use the individual portfolio shares at each age (see Figure 6) to weight the average return on each asset category reported in (Fagereng et al. 2020, Table 3). Given individual returns we compute within-group averages at each age.

in Norway, as reported in Fagereng, Holm, and Natvik (2021). In both scenarios, we assume that individuals fully reinvest their returns, following Fagereng et al. (2019).

Figure 11 displays the counterfactual wealth trajectories we compute. Portfolio and income heterogeneity are not sufficient to rationalize the patterns we observe for wealth accumulation. Instead, we would additionally require large and persistent differences in savings rates across the high-ranked and the middle-rise, on the one hand, and the low-ranked and the fallers, on the other. However, this would require 40 percent of the population to maintain an *average saving rate* of over 40 percent, as high as the *marginal propensity to save* out of unearned income, which is at odds with the evidence in Fagereng et al. (2019) for the saving behavior of Norwegians throughout the wealth distribution. Taken together, our results imply that persistent differences in returns (beyond those driven by portfolio composition) are needed in conjunction with differences in saving rates and income trajectories, as no single element can generate the joint patterns of income and wealth we observe.

FIGURE 11. Counterfactual wealth trajectories



Notes: The figure presents observed and counterfactual net worth profiles for the groups presented in Figure 3. We use the individual portfolio shares at each age (see Figure 6) to weight the average return on each asset category reported in (Fagereng et al. 2020, Table 3). Given these individual returns we compute the realized wealth using two hypothetical values for active saving out of (observed) non-capital income.

8. Parental background, education, and mobility

Having identified the importance of multidimensional heterogeneity for explaining wealth accumulation across the population, we now consider the role of individuals' ex-ante characteristics. Specifically, we determine whether the ex-post heterogeneity in life-cycle trajectories across groups can be predicted by factors observed early in life. We include characteristics emphasized in the intergenerational mobility literature, such as birthplace and parental wealth (Chetty et al. 2014; Boserup, Kopczuk, and

Kreiner 2017), and include education as another key determinant of labor income and intragenerational mobility.

8.1. Ex ante determinants of group membership

We quantify the predictive power of parental wealth, education, sex, and place of birth for group-assignment using a multinomial logit specification

$$\Pr\left(g_i = j\right) = F\left(\alpha_0^j + \beta_{q(i)}^j + \gamma_{\text{educ}(i)}^j + \delta_{\text{subj}(i)}^j + \lambda_{\text{pareduc}(i)}^j + \mu_{\text{male}(i)}^j + \nu_{\text{bplace}(i)}^j\right), \quad (6)$$

where $F(\cdot)$ denotes the logit transformation. Specifically, we include ventiles of parental wealth fixed effects, $\beta_{q(i)}^j$.[43] We also include education fixed effects (post-compulsory high school, technical college, undergraduate, post-graduate, and doctoral degree), $\gamma_{\text{educ}(i)}^j$, and subject-specific fixed effects, $\delta_{\text{subj}(i)}^j$, for individuals with undergraduate or graduate degrees.[44] Finally, we include a fixed effect for the highest level of parental education, $\lambda_{\text{pareduc}(i)}^j$, a sex fixed effect, $\mu_{\text{male}(i)}^j$, and birthplace fixed effects, $\nu_{\text{bplace}(i)}^j$, which allow for place-based differences among the Oslo metropolitan area, other major cities, and rural regions.

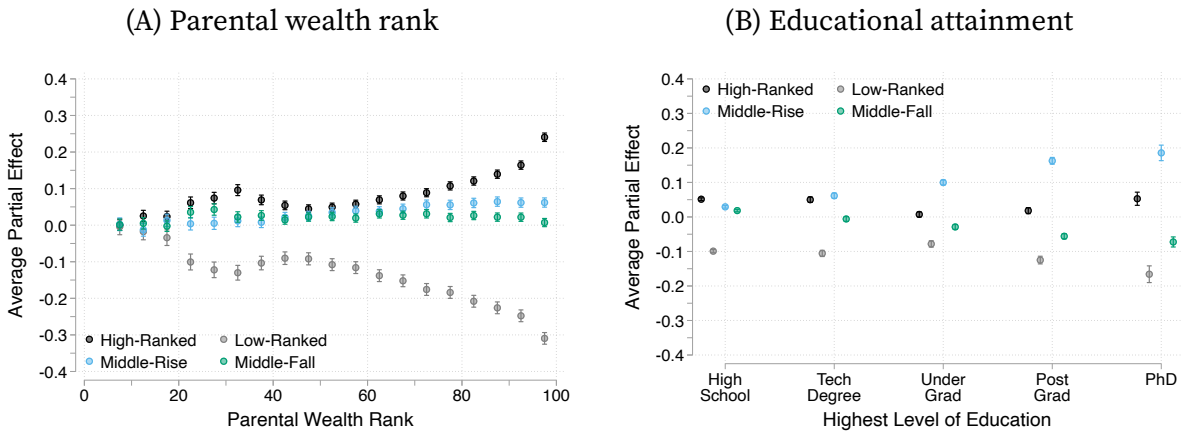
We find that parental wealth and the individuals' own education play a significant role in influencing group membership. These variables account for over 80 percent of the fit of the model in classifying individuals across groups. However, we also find that most of the uncertainty over group membership is not explained by these factors.[45]

[43] Formally, $q(i)$ is the ventile of the richest parent of individual i in the parent's own cohort wealth distribution in 1993 at the start of our sample.

[44] In practice, we aggregate the 350 degree-specific codes into six categories: arts and humanities; business, economics, and agricultural management; computer science and engineering; natural sciences; health; and education specialists.

[45] In alternative specifications we add covariates informative of parental background, such as parents' business ownership, and of the individuals' initial conditions, namely homeownership and business ownership and their initial wealth ventile. The average partial effects of the variables in (6) decrease 25 to 40 percent when conditioning on these additional variables, and the variation explained increases four-fold. These effects are driven by the individuals' initial wealth rank, consistent with the segmentation pattern we describe in Section 5.

FIGURE 12. Parental wealth, education and the probability of group assignment



Notes: Panel A plots the average partial effect of Parental Wealth (measured in 1993) relative to being born to parents in the bottom ventile of the distribution. Panel B plots the average partial effect of educational attainment relative to compulsory schooling age. We construct the average partial effect by integrating over the empirical joint distribution of other covariates. We report the probability of being assigned to each of our four groups, along with their 95 percent confidence intervals.

This finding is to be expected because part of the value of our clustering approach, which exploits the entire life-cycle history of individuals, lies in revealing a low-dimensional representation that captures important features of the data that are not easily summarized by observable variables.[46]

Parental wealth. Figure 12A reports the average partial effects of parental wealth rank on predicted group assignment and their 95 percent confidence intervals. The role of parental background is much stronger for the high- and low-ranked groups relative to the middle groups of risers and fallers. Individuals with wealthier parents are progressively more likely to belong to the high-ranked group and less likely to belong to the low-ranked group. An individual with parents in the top ventile of the wealth distribution is 25 percentage points more likely to belong to the high-ranked

[46]We calculate the share of variation in group membership explained by (6) and decompose the partial contribution of parental wealth, education, and initial characteristics using the Shapley-Owen decomposition (Shorrocks 2013) in Appendix E. This additive decomposition provides a single value per covariate category that is permutation invariant. We formally describe our methodology in Appendix F.

group than an individual with parents in the bottom ventile. The opposite is true for membership in the low-ranked group. By contrast, individuals coming from various parental backgrounds are almost equally likely to belong to the groups of risers and fallers up to parents in the top quartile of the wealth distribution.

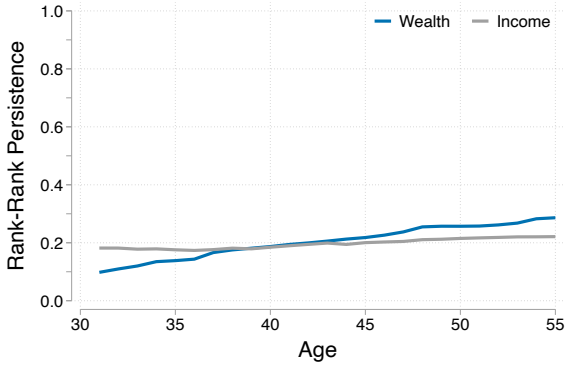
Education. Figure 12B reports the estimated average partial effects of educational attainment, along with their 95 percent confidence intervals. Higher educational attainment is associated with a higher probability of belonging to the group of risers, a much lower probability of belonging to the low-ranked, a lower probability of being a faller, and has a limited effect on the probability of being in the high-ranked group. Specifically, a university graduate is 10 percentage points more likely to belong to the risers' group than an individual with only compulsory schooling. This difference increases to up to 20 percentage points for PhDs, and provides evidence that risers are characterized by higher human capital that translates into increasing earnings profiles and wealth accumulation, as documented in Section 6. More generally, more educated individuals are less likely to be low-ranked. A postgraduate is 13 percentage points less likely to belong to the low-ranked than an individual with compulsory education.

Individuals of the high-ranked group are more likely to have a business or STEM degree, and those of the low-ranked group are less likely. Interestingly, the field of study is not relevant for distinguishing risers and fallers. Similarly, there is no significant role for parental education after taking into account the direct role of parental wealth and the individuals' own education as we show in Figure E.1 of Appendix E.

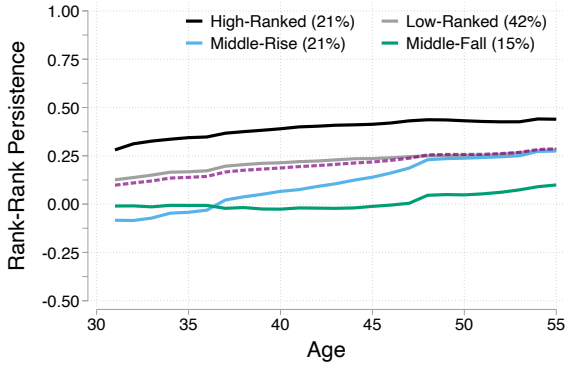
Sex and birthplace. We find a limited role for sex and birthplace conditional on parental wealth and education. Men are more likely to be in the high-ranked or falling group and correspondingly less likely to be risers or low-ranked. Birthplace has a small (but significant) effect on the probability of group membership. Being born in one of

FIGURE 13. Intergenerational persistence of ranks

(A) 1960–1964 birth cohort



(B) Four main groups



Notes: Panel A shows the intergenerational rank-rank persistence coefficient in wealth and income for the whole cohort. Panel B plots the intergenerational rank-rank wealth coefficient for the four main groups with the pooled cohort-level persistence measure shown with a dashed line.

the four largest Norwegian cities increases the probability of being in the high-ranked or rising groups by close to 5 percentage points. See Appendix E.

8.2. Relationship to intergenerational mobility

We document a trend of *declining intergenerational mobility* along the life cycle, which mirrors the increasing trend of intragenerational mobility in Section 3, and show how it relates to the wealth trajectories of the main mobility groups. We measure intergenerational mobility in wealth ranks as the correlation between individuals' rank and their parents' rank in 1993, when they were 55 years old on average.[47] Panel A of Figure 13 shows the estimates for the entire cohort. For comparison, we also report similar estimates for the intergenerational mobility of income ranks.

Intergenerational wealth mobility decreases with age, with intergenerational rank persistence increasing from 0.10 to 0.25, while income mobility remains the same as the

[47] Similarly to the definition in Equation (6), their parents' wealth rank is defined as that of the parent with the highest rank in their cohort's wealth distribution in 1993 at the start of our sample.

cohort grows older, at around 0.20.[48] The magnitude and increasing trend of the parent-child correlation in wealth ranks are similar to those reported by Boserup, Kopczuk, and Kreiner (2017) for Denmark, and slightly lower than the magnitudes reported by Adermon, Lindahl, and Waldenström (2018) and Black, Devereux, Lundborg, and Majlesi (2019) for Sweden.[49]

Panel B of Figure 13 shows that, despite significant level differences, the trend of declining intergenerational wealth mobility is common to all the main mobility groups, and particularly strong for the risers. The high-ranked are the least mobile group (the correlation with their parents' ranks increases from 0.30 to 0.45). At the other end, the fallers exhibit a much lower correlation with their parents' wealth rank throughout their lifetime (ranging from 0.00 to 0.10). Notably, the risers exhibit increasingly more intergenerational wealth persistence as they age (correlation going up to 0.25 starting from -0.10), which suggests it is more accurate to measure wealth ranks at the end of individuals' careers since some of the correlation with their parents' wealth materializes at that stage. Overall, these results are consistent with the role of parental wealth for individuals' intragenerational wealth mobility documented in Figure 12A.[50]

9. Robustness and method discussion

In this final section, we discuss the sensitivity of our results to a number of alternative methodological choices. Reassuringly, we find that our results do not substantively depend on either the clustering algorithm we choose to implement or the transformation of wealth we use to capture individuals' positions in the wealth distribution. Rather, we

[48] This flat pattern of intergenerational income mobility follows from the lower intragenerational mobility of income relative to wealth, which we document in Sections 3 and 7.

[49] Charles and Hurst (2003) report a coefficient of 0.37 for the correlation of log-wealth between generations in the US.

[50] Interestingly, although the increase in intergenerational wealth persistence is largest for the group rising through the wealth distribution, parental wealth is not a first-order determinant for this group once we control for education. This suggests that, for them, the role of parental wealth is mediated by investments in human capital.

find similar qualitative patterns of segmented wealth mobility for alternative clustering algorithms and outcomes, including measures of *absolute* mobility based on trajectories of (log) wealth levels. We also revisit the choice over the number of groups used in our baseline analysis by exploring the heterogeneity within each of the main four groups. Finally, we repeat our analysis for the 1965–69 cohort. We compare our findings and discuss the role of time effects and cross-cohort differences.

9.1. Robustness to alternative clustering algorithm and clustering variable

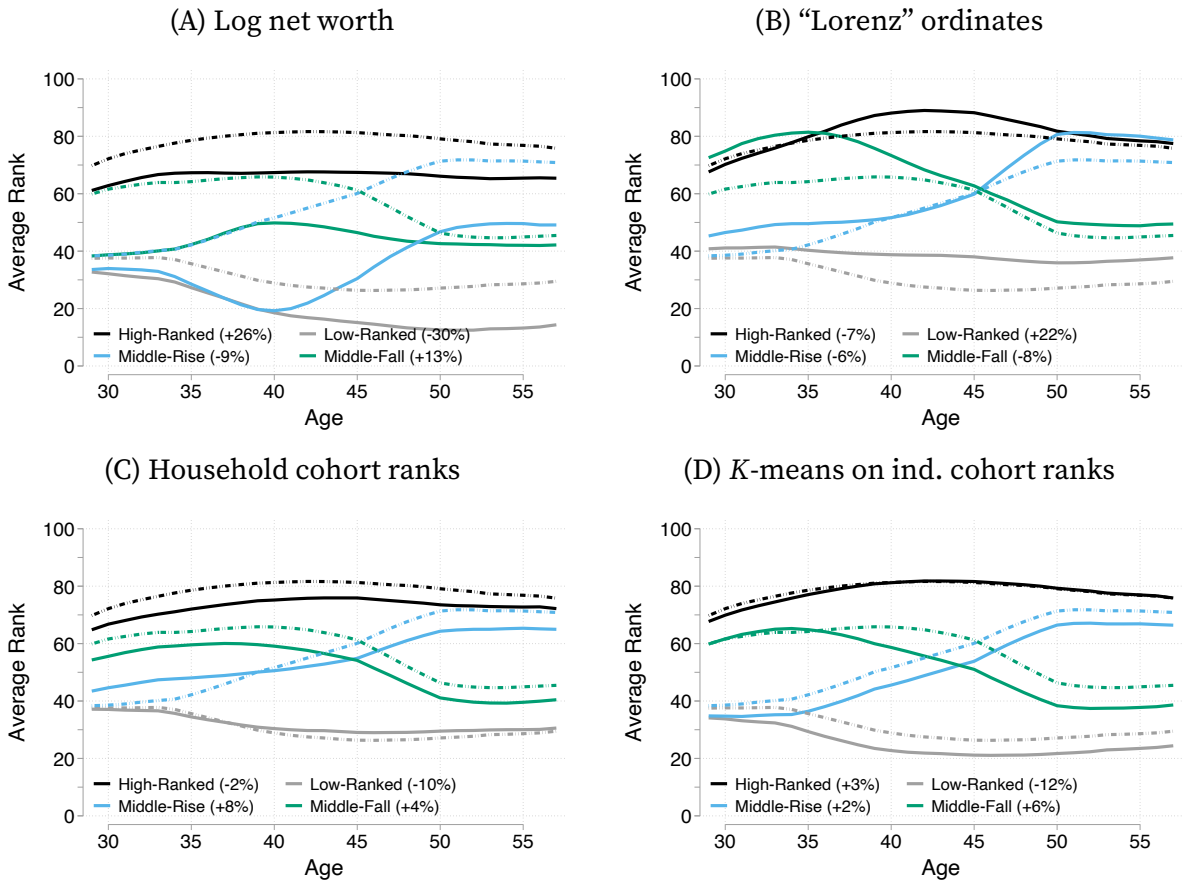
The qualitative pattern of segmented mobility that we document is a robust feature of the data. As we show in Figure 14, the mobility patterns we document emerge when using transformations of wealth histories that place more weight on differences in wealth levels, or when using household wealth ranks that take into account the role of marriage and cohabitation.[51] Similarly, our results are preserved when using the K-means algorithm to cluster individuals.

We begin by performing two robustness exercises in which we group individuals based on their log-wealth histories and the trajectories of their "Lorenz" ordinates using agglomerative hierarchical clustering, as in our baseline.[52] The log-wealth is a concave transformation of the wealth levels and magnifies differences for individuals with low levels of wealth, relative to those with high levels. Crucially, it provides a measure of *absolute* wealth mobility that complements the results on *relative* mobility we discussed in Sections 3 and 5. The Lorenz ordinates, given by individuals' positions in the cohort's Lorenz curve, measure the distance between individuals by the share of total wealth

[51] It is also not a mechanical result of clustering (as we already saw in Figure 9A when clustering on income trajectories). Simulation exercises confirm that, as the number of time periods increases, clustering recovers differences in the data-generating process across groups. This finding further highlights the importance of having a long panel of wealth histories to capture the typical mobility patterns of the population, as we mentioned in footnote 24.

[52] In both cases we bottom code wealth at 1,000 Norwegian Kroners to deal with negative values. Doing this has the added effect of compressing differences at the bottom of the wealth distribution.

FIGURE 14. Robustness of wealth rank clusters with $G = 4$



Notes: The figures plot the average wealth rank in each clustered group against the cohort's average age for clustering targeting different objectives. The clusters are constructed using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3). Panel D constructs clusters using the K-means algorithm. The dashed lines show the average wealth rank of the baseline clustering done with respect to individual cohort ranks, presented in Figure 3.

that lies between them.[53] Formally,

$$y_{i,t}^{Lorenz} = \frac{\sum_{\{j:w_{j,t} \leq w_{i,t}\}} w_{j,t}}{\sum_j w_{j,t}}. \tag{7}$$

This is a convex transformation of wealth ranks, and hence magnifies differences between individuals at the top of the distribution, making inequality more salient than

[53]The Lorenz curve maps each (ranked) individual to the cumulative share of wealth held by individuals who are poorer than they are.

either wealth ranks or log-wealth.

As Figures 14A and 14B show, these alternative clustering exercises do not change the defining mobility patterns described in Figure 3.[54] Nevertheless, the groups' composition and size do change as a result of the characteristics highlighted by the respective outcome variables. Using a concave transformation of wealth like the natural logarithm makes the high-ranked group larger and the low-ranked group smaller. By contrast, the Lorenz curve transformation makes the low-ranked group much larger as high wealth inequality makes the Lorenz curve flat at the bottom. Nevertheless, the patterns of the trajectories of each group are preserved.

In Figure 14C, we group individuals based on their household's wealth rank. This turns out to imply little changes for the resulting main groups. The main difference is a slight compression in the gaps between groups. The changes are ultimately small, in part because the Norwegian tax registry already treats couples by dividing equally the value of any joint assets, so there is little effect among married and cohabiting individuals who make up about 80 percent of the sample throughout (see Figure 8A).

Finally, the exercise shown in Figure 14D groups individual wealth rank trajectories with the K-means clustering algorithm. We set a number of $K = 4$ groups for comparison with our main exercise. The results are essentially the same as in our baseline, up to some minor recomposition across groups. For example, the low-ranked group is smaller and, as a result, poorer on average, while some of its members are reclassified as fallers, lowering that group's wealth rank profile.

9.2. Heterogeneity within groups

We now turn to examining the heterogeneity within each of the main four groups. We do so by exploiting the hierarchical nature of the clustering, that allows us to directly

[54] All the panels of Figure 14 report the typical individual wealth rank trajectories implied by the respective clustering exercise, regardless of the outcome variable used in the construction of the groups.

study the subgroups that make up each of the four major groups. We proceed by looking at the three main subgroups of each baseline group (see Figure C.1B for details).

We find further evidence of segmented wealth mobility within group by examining the typical trajectories of subgroups. We report each subgroup's typical wealth rank trajectory in Figure 15, along with that of the respective baseline group in light pink. Overall, mobility within each main group remains contained within segments of the wealth distribution, although there is overlap between groups, as shown by the interquartile range in Figure 3.

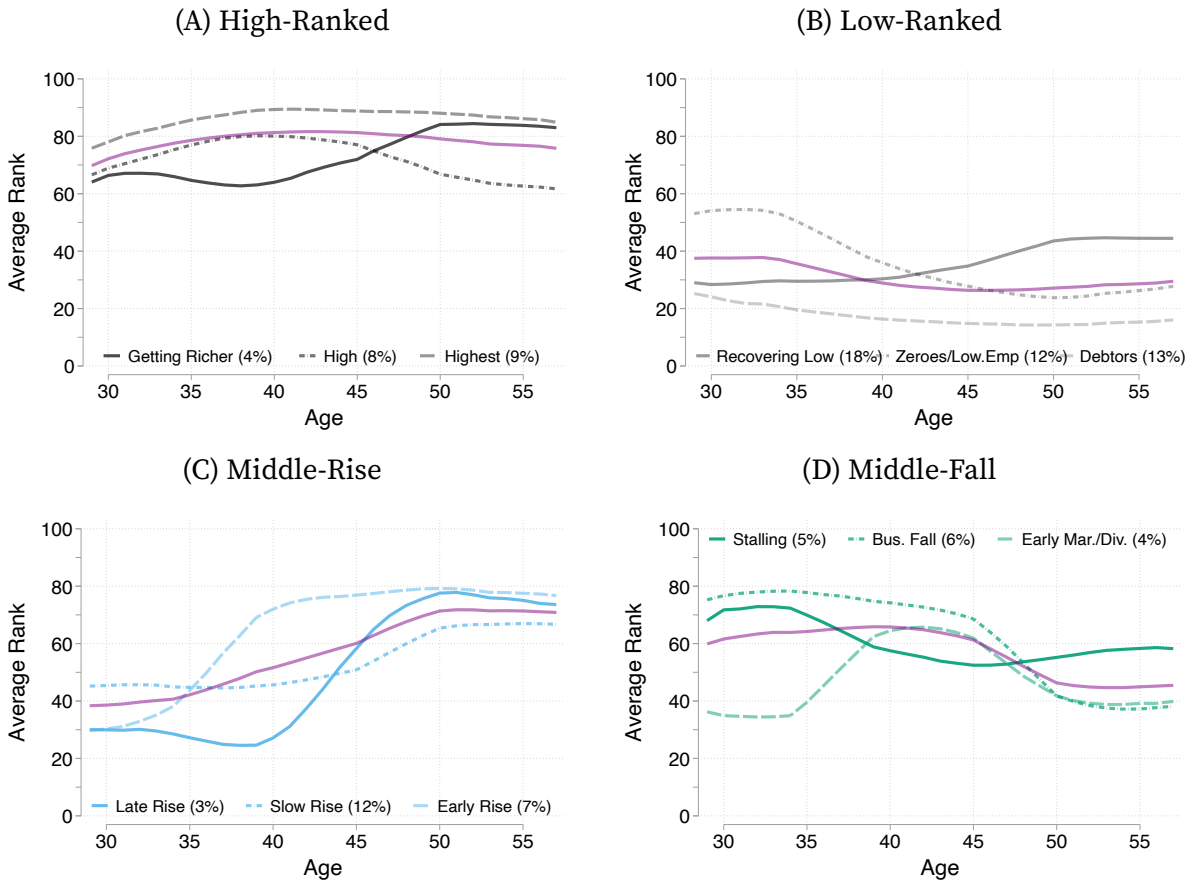
We now turn to describing the wealth mobility patterns that emerge for each subgroup. We examine their characteristics further in Appendix C.4.

High-ranked subgroups. Within the high-ranked, we find a group of individuals who are consistently at the top of the wealth distribution and two other groups that swap relative positions over the life cycle.[55] As with the baseline groups, we also find the subgroups differ in more than their wealth levels. Individuals in the group at the top have a consistently higher share of their wealth concentrated in privately held assets (such as businesses) and earn a larger fraction of their income from capital (including dividends), relative to individuals in the other two subgroups. By contrast, individuals in the rising subgroup, labeled as "getting richer" in Figure 15A, have a larger fraction of their wealth in property (although homeownership rates are similar between subgroups) and labor income makes up for the majority of their income. These characteristics make this group more similar to the main group of risers than the rest of the high-ranked group.

Low-ranked subgroups. The three subgroups at the bottom move similarly to those at the top. There is a subgroup of individuals that stay at the bottom of the wealth

[55]The groups and subgroups produced by our empirical strategy are designed to capture typical mobility patterns. The resulting typical trajectories turn out to capture the behavior of relatively large groups of individuals. For instance, the wealthiest group makes up 9 percent of our sample. See Ozkan et al. (2023) for a detailed discussion of individuals in the top 0.1 percent of the Norwegian wealth distribution.

FIGURE 15. Rank paths by subgroups



Notes: The figure plots the average wealth rank age profiles in the three main sub-clusters of each of the baseline clustered groups. Each panel corresponds to one of the four baseline groups reported in Figure 3. The light pink solid lines correspond to the average wealth rank in the corresponding main group. All individuals belong to the 1960–64 birth cohort. The clusters are constructed from the balanced sample using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3). The groups and subgroups are identified out of the dendrogram reported in Figure C.1B.

distribution throughout their lives and two groups that swap relative positions. We find that, throughout their lives, those lower in the distribution are net debtors; those falling always hold zero wealth, which makes them relatively poorer as the cohort ages; those recovering in ranks accumulate wealth after age 45, mostly in the form of property.

Middle-rise subgroups. These subgroups reveal differences in the timing of movements, with early and late risers that nevertheless begin and end the sample in

similar positions relative to each other (Figure 15C). Not surprisingly, late risers are more likely to earn graduate degrees and also take longer to acquire property, particularly relative to the “slow rise” subgroup.

***Middle-fall subgroups*.** Finally, the three subgroups within this group capture individuals who fall continuously through the distribution, who could not sustain their rise, and who stalled, remaining around the median of the wealth distribution as they grew older. The largest subgroup is characterized by a larger share of business owners, whose businesses produce a declining share of their income as they age, leading us to label this group "business fall".[56] The smaller subgroup has individuals who rise early but then fall after age 45. These individuals marry younger, on average, and have higher divorce rates relative to other groups; most of the increase in their wealth in their late 30s is tied to a rise in their homeownership rate.

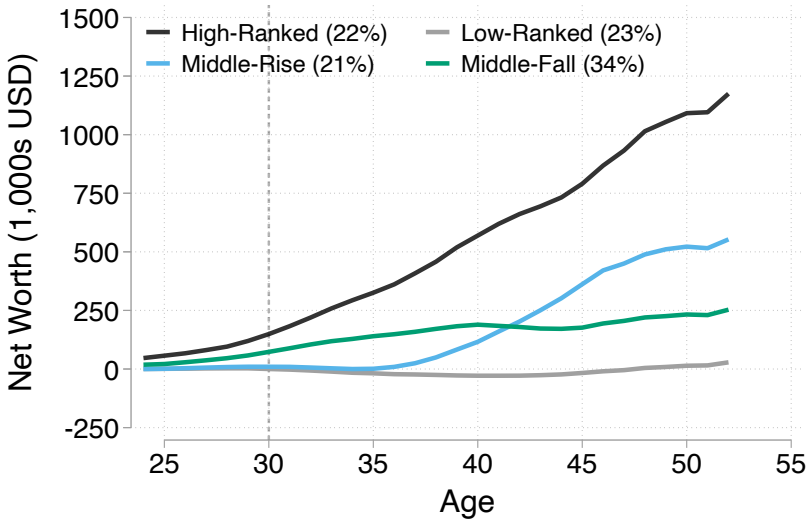
9.3. Time and age effects

Our results so far, based on the 1960–64 birth cohort, describe life-cycle patterns that encompass both age and calendar-time effects. However, the overall patterns of segmented mobility and wealth accumulation that we find are not exclusive to this cohort, even though the timing of aggregate events might affect groups differently in other cohorts. We show this for the 1965–69 birth cohort by replicating our clustering exercise using the latter 20 years of the sample, corresponding to the time they are 30 years or older, the same age as when the 1960–1964 cohort is first observed in the data. Figure 16 shows the results.

The general patterns are very similar to those documented for the main cohort, with two groups of persistently poor and rich individuals and two crossing groups in the

[56]The continuous decline of this group is accentuated by a substantial drop between ages 45 and 50. This drop coincides with the timing of the 2008 global financial crisis, hinting at a potential larger exposure of their businesses to foreign financial conditions. We turn to the role of time effects next.

FIGURE 16. Life-cycle dynamics of wealth mobility 1965 birth-cohort



Notes: The figure plots the average wealth level in each clustered group against individuals' age. All individuals belong to the 1965–69 birth cohort. The clusters are constructed using hierarchical agglomerative clustering and Ward's method with a dissimilarity measure (3). We use only the latter 20 years of data corresponding to the time for which the cohort is on average 30 years or older. These years are indicated in the figure by a vertical dashed line.

middle. Remarkably, the wealth levels for each group at each age are very similar to those of the 1960-64 cohort, although the wealth of the fallers stagnates slightly earlier, which could be explained by the fact that this cohort experienced the Great Recession at a younger age. Data availability does not allow us to repeat the main exercise for very distinct cohorts and, thus, completely separate age and time effects.

9.4. Alternative clustering approaches

Several alternative procedures have been proposed to construct latent groups from a sequence of realized outcomes. We briefly contrast our approach to two alternatives with applications to economic data.

First, the agglomerative hierarchical clustering method we employ is closely related to applications of Sequence Analysis tools that also approximate individual histories by

membership in a lower-dimensional group.[57] These tools are designed to summarize paths (or sequences) of categorical outcomes. Our approach is better suited to the continuous variation in wealth ranks that we consider because (*i*) it avoids categorizing ranks into arbitrary discrete groups and (*ii*) it lets us exploit the cardinality of our rank measurements. Accordingly, our approach captures both the ordinal and cardinal information in wealth ranks.

K-means clustering is another popular alternative procedure to retrieve latent groups from wealth rank histories (see Bonhomme and Manresa 2015; and Bonhomme, Lamadon, and Manresa 2022 for a discussion of the method and the derivation of its asymptotic distribution). Conceptually, *K*-means clustering uses an alternative distance metric to the one in equation (3), resulting in potentially distinct groups. We see hierarchical clustering as offering two important advantages over *K*-means clustering for large administrative datasets such as the one we analyze in this paper. First, *K*-means clustering involves solving an optimization problem through local optimization techniques that require many multi-start evaluations of the objective function for a given number of groups *G*. Although the solution is typically fast for a given initial guess of the partition, the computational demands become substantial for a large number of observations and when the dimensionality of each observation increases, as is our case.[58] Second, the procedure must be repeated whenever the number of target groups *G* changes. By contrast, once the dissimilarity measure is specified, hierarchical clustering recovers all optimal clusters for $G \in \{1, \ldots, N\}$ as an outcome of a global search for the sequential agglomeration of groups.

Our clustering procedure, Sequence Analysis, and *K*-means all generate partitions of the set of individuals into disjoint groups. An alternative approach is to generate

[57] These tools originate in quantitative sociology (Dijkstra and Taris 1995; McVicar and Anyadike-Danes 2002; Dlouhy and Biemann 2015) and have been applied in economics (for example, Humphries 2022).

[58] Given the data security requirements for using datasets such as ours, there are limitations to the computational resources we can access.

probabilistic assignments of individuals to groups. For instance, Lewis, Melcangi, and Pilossoph (2021) use a variation of the K-means algorithm to recover "fuzzy clusters" from cross-sectional data. Relatedly, Ahn, Hobijn, and Şahin (2023) use a hidden Markov chain model to identify latent types of workers and probabilistically assign observed workers to these types using the EM algorithm.

10. Conclusion and directions for future research

In this paper, we used 25 years of administrative records on the wealth, income, and other characteristics of Norwegians to study the life-cycle dynamics of wealth mobility. Intragenerational mobility increases throughout the life cycle, but most of it is driven by two broad groups of individuals in the middle of the distribution that switch position as they age. The flipside of this result is that wealth mobility is limited for many individuals: 42 percent remain near the bottom of the wealth distribution over their work life, while 21 percent are always near the top. We identified these patterns by studying the dynamics of wealth accumulation for a full cohort of Norwegian residents, grouping them based on their realized trajectories across the wealth distribution.

Our approach is uniquely suited to understand how frequent reversals of fortune are, when they occur, and how persistent they are. Importantly, a snapshot of the wealth distribution or a measure of wealth mobility computed on a shorter panel would not contain enough information to recognize the differential paths of the risers and fallers, or to establish that the high-ranked tend to only fluctuate around the upper part of the wealth distribution throughout their lives. In this sense, our results complement the extensive literature on wealth inequality and top wealth shares (Saez and Zucman 2016; Smith, Zidar, and Zwick 2022) and recent work on the dynamics of wealth for the wealthiest (Gomez 2023; Ozkan, Hubmer, Salgado, and Halvorsen 2023).

We also showed that family background and education are key determinants of

life-cycle wealth paths, but many other observable and unobservable factors are also at play. The groups we recovered display distinct profiles in terms of labor market income, portfolio decisions, and business ownership. This exemplifies how our clustering method provides a simple way of establishing relationships among these many different variables and an individual's wealth path. Furthermore, the joint patterns of wealth and income mobility we document are at odds with standard homothetic life-cycle savings models, implying a necessary role for persistent heterogeneity in saving rates and returns to explain wealth dynamics for individuals throughout the distribution and not just for those at the very top.

References

Adermon, Adrian, Mikael Lindahl, and Daniel Waldenström. 2018. "Intergenerational wealth mobility and the role of inheritance: Evidence from multiple generations." *The Economic Journal* 128 (612):F482–F513. URL https://doi.org/10.1111/ecoj.12535.

Ahn, Hie Joo, Bart Hobijn, and Ayşegül Şahin. 2023. "The dual U.S. labor market uncovered." Finance and Economics Discussion Series 2023–031, Board of Governors of the Federal Reserve System. URL https://doi.org/10.17016/FEDS.2023.031.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. "Earnings and consumption dynamics: A nonlinear panel data framework." *Econometrica* 85 (3):693–734. URL https://doi.org/10.3982/ECTA13795.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6):2333—-67. URL https://doi.org/10.1257/aer.20160137.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2017. "Earnings Inequality and Other Determinants of Wealth Inequality." *American Economic Review* 107 (5):593–97. URL https://doi.org/10.1257/aer.p20171005.

Bhuller, Manudeep, Magne Mogstad, and Kjell G Salvanes. 2017. "Life-cycle earnings, education premiums, and internal rates of return." *Journal of Labor Economics* 35 (4):993–1030. URL https://doi.org/10.1086/692509.

Black, Sandra E, Paul J Devereux, Fanny Landaud, and Kjell G Salvanes. 2022. "The (un) importance of inheritance." Tech. rep., National Bureau of Economic Research. URL https://doi.org/10.3386/w29693.

———. 2023. "Where does wealth come from? Measuring lifetime resources in Norway." *Journal of Economic Perspectives* 37 (4):115–136. URL https://doi.org/10.1257/jep.37.4.115.

Black, Sandra E, Paul J Devereux, Petter Lundborg, and Kaveh Majlesi. 2019. "Poor Little Rich Kids? The Role of Nature versus Nurture in Wealth and Other Economic Outcomes and Behaviours." *The Review of Economic Studies* 87 (4):1683–1725. URL https://doi.org/10.1093/restud/rdz038.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes. 2005. "Why the apple doesn't fall far: Understanding intergenerational transmission of human capital." *American Economic Review* 95 (1):437–449. URL https://doi.org/10.1257/0002828053828635.

Blau, David M. 2016. "Pensions, household saving, and welfare: A dynamic analysis of crowd out." *Quantitative Economics* 7 (1):193–224. URL https://doi.org/10.3982/QE349.

Blundell, Richard, Michael Graber, and Magne Mogstad. 2015. "Labor income dynamics and the insurance from taxes, transfers, and the family." *Journal of Public Economics* 127:58–73. URL https://doi.org/10.1016/j.jpubeco.2014.04.011.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A distributional framework for matched employer employee data." *Econometrica* 87 (3):699–739. URL https://doi.org/10.3982/ECTA15722.

———. 2022. "Discretizing unobserved heterogeneity." *Econometrica* 90 (2):625–643. URL https://doi.org/10.3982/ECTA15238.

Bonhomme, Stéphane and Elena Manresa. 2015. "Grouped patterns of heterogeneity in panel

data." *Econometrica* 83 (3):1147–1184. URL https://doi.org/10.3982/ECTA11319.

Borysov, Petro, Jan Hannig, and J.S. Marron. 2014. "Asymptotics of hierarchical clustering for growing dimension." *Journal of Multivariate Analysis* 124:465–479. URL https://doi.org/10.1016/j.jmva.2013.11.010.

Boserup, Simon Halphen, Wojciech Kopczuk, and Claus Thustrup Kreiner. 2017. "Intergenerational wealth formation over the life cycle: Evidence from Danish wealth records 1984-2013." Working paper, University of Copenhagen. URL https://www.columbia.edu/~wk2110/bin/WealthAcrossGen.pdf.

———. 2018. "Born with a silver spoon? Danish evidence on wealth inequality in childhood." *The Economic Journal* 128 (612):F514–F544. URL https://doi.org/10.1111/ecoj.12496.

Bénabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica* 70 (2):481–517. URL https://doi.org/10.1111/1468-0262.00293.

Carroll, Christopher D. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 23 (2):61–156. URL https://www.brookings.edu/wp-content/uploads/1992/06/1992b_bpea_carroll_hall_zeldes.pdf.

Charles, Kerwin Kofi and Erik Hurst. 2003. "The correlation of wealth across generations." *Journal of Political Economy* 111 (6):1155–1182. URL https://doi.org/10.1086/378526.

Chetty, Raj, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang. 2017. "The fading American dream: Trends in absolute income mobility since 1940." *Science* 356 (6336):398–406. URL https://doi.org/10.1126/science.aal4617.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. 2014. "Is the United States S¡still a land of opportunity? Recent trends in intergenerational mobility." *American Economic Review* 104 (5):141–47. URL https://doi.org/10.1257/aer.104.5.141.

Chetverikov, Denis and Daniel Wilhelm. 2023. "Inference for rank-rank regressions." econ.EM preprint 2310.15512, arXiv. URL https://doi.org/10.48550/arXiv.2310.15512.

Davies, James B. and Anthony F. Shorrocks. 2000. "The distribution of wealth." In *Handbook of Income Distribution*, vol. 1, edited by Anthony B. Atkinson and François Bourguignon, chap. 11. Elsevier, 605–675. URL https://doi.org/10.1016/S1574-0056(00)80014-7.

De Nardi, Mariacristina and Giulio Fella. 2017. "Saving and wealth inequality." *Review of Economic Dynamics* 26:280–300. URL https://doi.org/10.1016/j.red.2017.06.002.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2020. "Nonlinear household earnings dynamics, self-insurance, and welfare." *Journal of the European Economic Association* 18 (2):890–926. URL https://doi.org/10.1093/jeea/jvz010.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5):1221. URL https://doi.org/10.2307/2938366.

Deutscher, Nathan and Bhashkar Mazumder. 2023. "Measuring intergenerational income mobility: a synthesis of approaches." *Journal of Economic Literature* 61 (3):988–1036. URL https://doi.org/10.1257/jel.20211413.

Dijkstra, Wil and Toon Taris. 1995. "Measuring the agreement between sequences." *Sociological Methods & Research* 24 (2):214–231. URL https://doi.org/10.1177/0049124195024002004.

Dlouhy, Katja and Torsten Biemann. 2015. "Optimal matching analysis in career research: A

review and some best-practice recommendations." *Journal of Vocational Behavior* 90:163–173. URL https://doi.org/10.1016/j.jvb.2015.04.005.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the rich save more?" *Journal of Political Economy* 112 (2):397–444. URL https://doi.org/10.1086/381475.

Egashira, Kento, Kazuyoshi Yata, and Makoto Aoshima. 2024. "Asymptotic properties of hierarchical clustering in high-dimensional settings." *Journal of Multivariate Analysis* 199:105251. URL https://doi.org/10.1016/j.jmva.2023.105251.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and persistence in returns to wealth." *Econometrica* 88 (1):115–170. URL https://doi.org/10.3982/ECTA14835.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2022. "Assortative mating and wealth inequality." Working Paper 29903, National Bureau of Economic Research. URL http://doi.org/10.3386/w29903.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC heterogeneity and household balance sheets." *American Economic Journal: Macroeconomics* 13 (4):1–54. URL https://doi.org/10.1257/mac.20190211.

Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. 2019. "Saving behavior across the wealth distribution: The importance of capital gains." Working Paper 26588, National Bureau of Economic Research. URL https://doi.org/10.3386/w26588.

Fagereng, Andreas, Martin Blomhoff Holm, and Kjersti Næss Torstensen. 2020. "Housing wealth in Norway, 1993–2015." *Journal of Economic and Social Measurement* 45 (1):65–81. URL https://doi.org/10.3233/JEM-200471.

Fagereng, Andreas, Magne Mogstad, and Marte Rønning. 2021. "Why do wealthy parents have wealthy children?" *Journal of Political Economy* 129 (3):703–756. URL https://doi.org/10.1086/712446.

Gale, William G., Hilary Gelfond, Jason J. Fichtner, and Benjamin H. Harris. 2021. "The Wealth of generations, with special attention to the millennials." In *Measuring Distribution and Mobility of Income and Wealth*, edited by Raj Chetty, John N. Friedman, Janet C. Gornick, Barry Johnson, and Arthur Kennickell. University of Chicago Press, 145–174. URL http://www.nber.org/chapters/c14445.

Gomez, Matthieu. 2023. "Decomposing the growth of top wealth shares." *Econometrica* 91 (3):979–1024. URL https://doi.org/10.3982/ECTA16755.

Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (1):47–89. URL https://doi.org/10.1111/1468-0262.00269.

Gregory, Victoria, Guido Menzio, and David G Wiczer. 2021. "The alpha beta bamma of the labor market." Working Paper 28663, National Bureau of Economic Research. URL https://doi.org/10.3386/w28663.

Guvenen, Fatih, Gueorgui Kambourov, Burhan Kuruscu, Sergio Ocampo, and Daphne Chen. 2023. "Use It or Lose It: Efficiency and Redistributional Effects of Wealth Taxation." *The Quarterly Journal of Economics* 138 (2):835–894. URL https://doi.org/10.1093/qje/qjac047.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2022. "Lifetime earnings in the United States over six decades." *American Economic Journal: Applied Economics* 14 (4):446–79.

URL https://doi.org/10.1257/app.20190489.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?" *Econometrica* 89 (5):2303–2339. URL https://doi.org/10.3982/ECTA14603.

Halvorsen, Elin, Serdar Ozkan, and Sergio Salgado. 2022. "Earnings dynamics and its intergenerational transmission: Evidence from Norway." *Quantitative Economics* 13 (4):1707–1746. URL https://doi.org/10.3982/QE1849.

Hastie, Trevor, Robert Tibshirani, and Jerome H Friedman. 2009. *The elements of statistical learning: Data mining, inference, and prediction*, vol. 2. New York, NY: Springer.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review* 104 (7):2075–2126. URL https://doi.org/10.1257/aer.104.7.2075.

Hebous, Shafik, Zhiyang Jia, Knut Løyland, Thor O Thoresen, and Arnstein Øvrum. 2023. "Do audits improve future tax compliance in the absence of penalties? Evidence from random audits in Norway." *Journal of Economic Behavior & Organization* 207:305–326. URL https://doi.org/10.1016/j.jebo.2023.01.001.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith. 2021. "Sources of US wealth inequality: Past, present, and future." *NBER Macroeconomics Annual* 35:391–455. URL https://doi.org/10.1086/712332.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of lifetime inequality." *American Economic Review* 101 (7):2923–2954. URL https://doi.org/10.1257/aer.101.7.2923.

Humphries, John Eric. 2022. "The causes and consequences of self-employment over the life cycle." URL https://johnerichumphries.com/Humphries_SelfEmployment.pdf. Working Paper.

Hurst, Erik, Ming Ching Luoh, Frank P. Stafford, and William G. Gale. 1998. "The wealth dynamics of American families, 1984-94." *Brookings Papers on Economic Activity* 1998 (1):267–337. URL https://doi.org/10.2307/2534673.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3):697-–743. URL https://doi.org/10.1257/aer.20160042.

Krueger, Alan B. 2012. "The rise and consequences of inequality in the United States." URL https://obamawhitehouse.archives.gov/sites/default/files/krueger_cap_speech_final_remarks.pdf.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2020. "Income and wealth inequality in America, 1949–2016." *Journal of Political Economy* 128 (9):3469–3519. URL https://doi.org/10.1086/708815.

Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph. 2021. "Latent heterogeneity in the marginal propensity to consume." Staff Reports 902, Federal Reserve Bank of New York. URL https://www.newyorkfed.org/research/staff_reports/sr902.html.

Low, Hamish and Luigi Pistaferri. 2015. "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off." *American Economic Review* 105 (10):2986—-3029. URL https://doi.org/10.1257/aer.20110108.

McVicar, Duncan and Michael Anyadike-Danes. 2002. "Predicting successful and unsuccessful

transitions from school to work by using sequence methods." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 165 (2):317–334. URL https://doi.org/10.1111/1467-985X.00641.

Meghir, Costas and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72 (1):1–32. URL https://doi.org/10.1111/j.1468-0262.2004.00476.x.

———. 2011. "Earnings, Consumption and Life Cycle Choices." In *Handbook of Labor Economics*, vol. 4, edited by David Card and Orley Ashenfelter. Elsevier, 773–854. URL https://doi.org/10.1016/S0169-7218(11)02407-5.

Ozkan, Serdar, Joachim Hubmer, Sergio Salgado, and Elin Halvorsen. 2023. "Why are the wealthiest so wealthy? A longitudinal empirical investigation." Working Paper 2023–004, Federal Reserve Bank of St. Louis. URL https://doi.org/10.20955/wp.2023.004.

Paz-Pardo, Gonzalo. 2024. "Homeownership and portfolio choice over the generations." *American Economic Journal: Macroeconomics* 16 (1):207–237. URL https://doi.org/10.1257/mac.20200473.

Phelan, Thomas. 2025. "On the Optimality of Differential Asset Taxation." *International Economic Review* :1–47URL https://doi.org/10.26509/frbc-wp-201917.

Sabelhaus, John. 2024. "Lifetime Wealth Transfers and Dynamic Wealth Accumulation." URL https://sites.google.com/prod/view/johnsabelhaus/home. Working Paper.

Saez, Emmanuel and Gabriel Zucman. 2016. "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data." *The Quarterly Journal of Economics* 131 (2):519–578. URL https://doi.org/10.1093/qje/qjw004.

Samwick, Andrew. 2003. "Patience, pensions, and saving." *Unpublished Manuscript, Darmouth College* URL https://www.researchgate.net/publication/228932485_Patience_Pensions_and_Saving.

Sargent, Thomas J, Neng Wang, and Jinqiang Yang. 2021. "Stochastic Earnings Growth and Equilibrium Wealth Distributions." Working Paper 28473, National Bureau of Economic Research. URL http://doi.org/10.3386/w28473.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving "Optimally" for Retirement?" *Journal of Political Economy* 114 (4):607–643. URL https://doi.org/10.1086/506335.

Shiro, Ariel Gelrud, Christopher Pulliam, John Sabelhaus, and Ember Smith. 2022. "Stuck on the ladder: Intragenerational wealth mobility in the United States." Working paper, Brookings Institution. URL https://www.brookings.edu/wp-content/uploads/2022/06/2022_FMCI_IntragenerationalWealthMobility_FINAL.pdf.

Shorrocks, Anthony F. 1978. "The measurement of mobility." *Econometrica* 46 (5):1013–1024. URL https://doi.org/10.2307/1911433.

———. 2013. "Decomposition procedures for distributional analysis: A unified framework based on the Shapley value." *Journal of Economic Inequality* 11 (1):99. URL https://doi.org/10.1007/s10888-011-9214-z.

Smith, Matthew, Owen Zidar, and Eric Zwick. 2022. "Top wealth in America: New estimates under heterogeneous returns." *The Quarterly Journal of Economics* 138 (1):515–573. URL https://doi.org/10.1093/qje/qjac033.

Solon, Gary. 1992. "Intergenerational income mobility in the United States." *The American*

Economic Review 82 (3):393–408. URL http://www.jstor.org/stable/2117312.

Stachurski, John and Alexis Akira Toda. 2019. "An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity." *Journal of Economic Theory* 182:1–24. URL https://doi.org/10.1016/j.jet.2019.04.001.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk." *Journal of Political Economy* 112 (3):695–717. URL https://doi.org/10.1086/383105.

Straub, Ludwig. 2019. "Consumption, savings, and the distribution of permanent income." *Unpublished manuscript, Harvard University* 17. URL https://scholar.harvard.edu/files/straub/files/cons_ineq_rates.pdf.

Thoresen, Thor O, Marius AK Ring, Odd E Nygård, and Jon Epland. 2022. "A wealth tax at work." *CESifo Economic Studies* 68 (4):321–361. URL https://doi.org/10.1093/cesifo/ifac009.

Vermeulen, Philip. 2016. "Estimating the top tail of the wealth distribution." *American Economic Review* 106 (5):646–50. URL https://doi.org/10.1257/aer.p20161021.

Wedgwood, Josiah. 1929. *The economics of inheritance*. Routledge.

Zeldes, Stephen P. 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *The Quarterly Journal of Economics* 104 (2):275–298. URL https://doi.org/10.2307/2937848.

Appendix A. Detailed description of the measurement of wealth

TABLE A.1. Data Sources and Variables

Source	Variable	Metadata
Befolkningsstatistikk		
Fixed variables	Individual annonimized ID	https://www.ssb.no/a/metadata/codelist/datadok/1668202/no
	Mother's, Father's, Spouse's, Cohabitant's, & Family ID	
	Sex	https://www.ssb.no/a/metadata/codelist/datadok/1385084/no
	Birth city	https://www.ssb.no/a/metadata/codelist/datadok/637597/no
	Immigration category	https://www.ssb.no/a/metadata/solr.cgi?q=invkat
	Birth date	
	Immigration: First year of record	
	Emmigration date	
Family and cohabitation	Civil status	https://www.ssb.no/en/klass/klassifikasjoner/19
Utdanningsstatistikk		
	Highest level of education	https://www.ssb.no/en/klass/klassifikasjoner/36
	Field of study	https://www.ssb.no/en/klass/klassifikasjoner/36
Inntekts og formuesstatistikk		
Selvangivelsesregisteret	Taxable net worth	https://www.ssb.no/a/metadata/conceptvariable/vardok/18/nb
(Section 4 –Formue– of the tax	Taxable assets	https://www.ssb.no/a/metadata/conceptvariable/vardok/662/nb
declaration form)	Debt	https://www.ssb.no/a/metadata/conceptvariable/vardok/17/nb
	Inheritance/Gifts	https://www.ssb.no/a/metadata/conceptvariable/vardok/3494/nb
	Bank deposits	https://www.ssb.no/a/metadata/conceptvariable/vardok/591/nb
	Bonds (VPS-registered)	https://www.ssb.no/a/metadata/conceptvariable/vardok/1210/nb
	Shares (VPS-registered)	https://www.ssb.no/a/metadata/conceptvariable/vardok/3109/nb
	Mutual funds	https://www.ssb.no/a/metadata/conceptvariable/vardok/679/nb
	Foreign deposits	https://www.ssb.no/a/metadata/conceptvariable/vardok/3239/nb
Housing wealth database	Value of housing (including cabins and secondary homes) (Adjusted using \citealt{fagereng2020housing})	
Inntektsstatistikk	Employment income (incl. par. & sickness ben.)	https://www.ssb.no/a/metadata/conceptvariable/vardok/15/en
	Self-emp. income (incl. sickness ben.)	https://www.ssb.no/a/metadata/conceptvariable/vardok/13/en
	Capital income	https://www.ssb.no/a/metadata/conceptvariable/vardok/10/en
	Interest income	https://www.ssb.no/a/metadata/conceptvariable/vardok/560/nb
	Dividends	https://www.ssb.no/a/metadata/conceptvariable/vardok/561/en
	Realized capital gains	https://www.ssb.no/a/metadata/conceptvariable/vardok/562/en
	Realized capital losses	https://www.ssb.no/a/metadata/conceptvariable/vardok/563/en
	Social security benefits	https://www.ssb.no/a/metadata/conceptvariable/vardok/615/en
	Retirement pension	https://www.ssb.no/a/metadata/conceptvariable/vardok/25/nb
	Disability pension	https://www.ssb.no/a/metadata/conceptvariable/vardok/33/nb
	Unemployment benefits	https://www.ssb.no/a/metadata/conceptvariable/vardok/666/en
	Sickness benefits	https://www.ssb.no/a/metadata/conceptvariable/vardok/3366/en
	Parental benefits	https://www.ssb.no/a/metadata/conceptvariable/vardok/3367/en
	Child allowance	https://www.ssb.no/a/metadata/conceptvariable/vardok/567/en
	Dwelling support	https://www.ssb.no/a/metadata/conceptvariable/vardok/667/en
	Student grant	https://www.ssb.no/a/metadata/conceptvariable/vardok/576/en
	Child care benefits	https://www.ssb.no/a/metadata/conceptvariable/vardok/568/en

Notes: The table reports selected variables used in the analysis along with their source and corresponding metadata. Not all variables are included for space. Variable definitions come from https://www.ssb.no/a/metadata/definisjoner/variabler/main.html.

Tax Returns. We construct our main variable of interest, a measure of an individual's wealth, using data from tax returns. We use tax returns for all individuals who are tax resident in Norway between 1993 and 2017. Norwegians who live abroad are not tax resident in Norway and do not have to file a tax return in that year. All values are measured on December 31st of each year. Thus, we observe an annual snapshot of an individual's balance sheet.

We observe both income tax records and wealth tax records. The income tax is levied on individuals regardless of marital status. In contrast, wealth is jointly taxed for married couples. Cohabitant couples (with or without children) are taxed separately even though they may own assets jointly. For jointly owned assets, the division of assets follows formal ownership share in the case of housing, but couples can elect to assign mortgages to individual tax records in a different proportion. The majority of individuals assets on the tax records are divided equally because only a small minority of individuals actually pay wealth taxes. This is because the wealth tax is levied on the value of 'taxable wealth' which differs from market wealth because the tax code includes discounts on a number of assets (Thoresen et al. 2022). We use the discounted values that appear on individuals' tax records to reverse engineer the tax assessed market value of their wealth.

Similarly, for income tax records we observe taxable incomes, taxes paid, and benefits received at the individual level. We are able to observe both total values and disaggregated values into broad categories. For instance, this allows us to observe employee income separately from self-employed income, interest income, or dividend income. We construct indicators for receiving a type of income or benefit by assigning the value of 1 to all individuals reporting a positive flow in that year and a value of 0 otherwise.

The majority of the information in these tax records is third-party reported. For example, employers report employee income directly to the tax authority and, similarly, financial institutions (banks, brokers, the Norwegian Central Securities Depository–VPS, etc.) report the value of assets held in their accounts. This greatly reduces concerns of measurement error (e.g., due to recall biases) or misreporting. However, some components of the tax return are not reported by third parties. These include foreign income or dividends not registered in the VPS and the values of some other assets—including foreign assets and valuables including art, paintings, and jewelry.[59] Individuals are responsible for disclosing the values of these incomes and assets. They are required to substantiate reported valuations (or deductions) which are checked by the tax authority (see Hebous et al. 2023 for a discussion of the Norwegian Tax Authority auditing process). In addition, the tax authority has automated control routines that flag tax returns with extreme movements in either income or wealth (which may indicate large evasion), and the tax authority checks these returns in more detail. This eliminates measurement error from self-reporting and censoring that is common in household surveys.[60]

[59] We include these last two wealth components when constructing the market value of wealth. However, when we disaggregated wealth and discuss portfolio components we do not report results for these categories. This is for two reasons. First, they are only important for a small share of individuals and, second, because are self reported and hard to value we view them as less reliable when disaggregated. As we discuss in the main text, the nature of our rank measure minimizes the impact of these concerns.

[60] Measurement error from censoring and misreporting has a limited effect on mean estimates, but represents an important challenge for the study of dynamics of individual observations, typically attenuating persistence measures. In the income dynamics literature this has lead to the popularity of errors in variables estimators.

Housing wealth database. The Norwegian tax authority directly estimates the market value of housing using a procedure that updates original transaction prices for houses that are not transacted. As Fagereng, Holm, and Torstensen (2020) discuss, this can lead to systematic measurement error in the valuation of housing wealth. We use an updated version of the correction series reported in Fagereng, Holm, and Torstensen (2020) for the period 1993 to 2017 provided by the authors to adjust the value of housing.[61] The imputation procedure applies a different correction to owner-occupied housing, secondary housing, and cabins (holiday homes). Finally, as we describe above, we allocate the fraction of the house owned by an individual to each individual's market wealth. We aggregate primary residences, secondary residences and leisure properties, and foreign residences into a property-asset class. Finally, we define a home ownership indicator excluding secondary and foreign properties.

Inheritance tax records. Information about inheritances and inter-vivos gifts as derived from Inheritance tax records 1995-2013. The inheritance tax in Norway was abolished in 2014. This is distinct from the annually levied wealth tax. The tax registry reports gifts and inheritances above 10,000 NOK between 1993 and 1996 and above 100,000 NOK after 1996. Absent the tax relevance of this information the values are not subject to verification and the number of individuals reporting gifts or inheritances is severely reduced after 2014.

Central population register and Norwegian educational database. This is an annual national register available since 1964. It contains anonymized individual identification numbers, residence, marital status, highest completed education, and field of study. From this we are able to link spouses, children to their parents, and calculate educational attainment. We additionally use this in our sample selection to construct an individual's migration history.

We have detailed information on individual education levels and fields of study, according to the Norwegian Standard Classification of Education (NUS2000).[62] This classification provides nine levels of education, ranging from no education to post-graduate PhD level, as well as 350 fields of study. The 9 levels of education corresponding to the NUS 1 digit categories are: No education or preschool education; Primary education; Lower secondary education; Upper secondary basic education; Upper education final year education; Post-secondary but non-tertiary education; Undergraduate degree; Graduate degree up to PhD(c); Post-graduate level (PhD). We aggregate the fields of study at the NUS 3 digit level into six groups: Humanities; Economics, Business, Law and Management; Computer Science and Engineering; Natural Sciences; Medicine and Health Care; and Teaching. To construct our education levels We pool all post-graduate degrees that do not result in the award of a doctorate

[61] Fagereng, Holm, and Torstensen (2020) use an ensemble machine learning method on housing transaction data to produce their correction series. They show that this has considerable in- and out-of-sample performance improvements over comparable hedonic price regression based imputations.

[62] See https://www.ssb.no/en/klass/klassifikasjoner/36.

and, additionally, pool all education levels less than high school. Finally, we use the group "technical degree" for those additional levels of education that do not result in an undergraduate degree, but require additional study beyond high school.

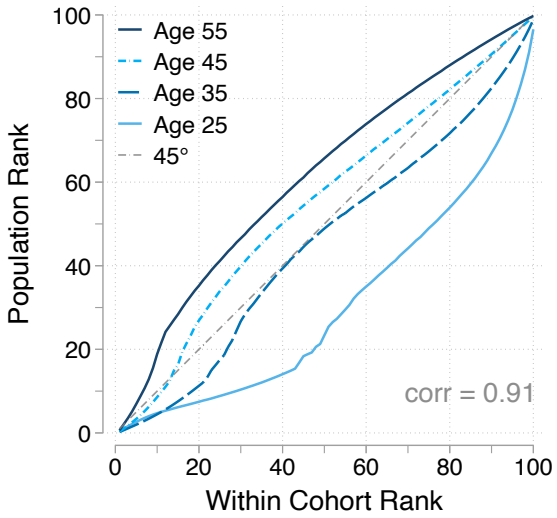
Appendix B. Additional results on measurement of wealth ranks

B.1. Cohort wealth ranks vs population wealth ranks

The dispersion of wealth within our cohort of interest increases as the cohort ages. This means that the correspondence between wealth ranks and wealth levels changes over time. Figure B.1 visualizes this by plotting the correspondence between within-cohort ranks and population ranks at different ages. The correspondence between population wealth ranks and wealth levels is shown in Figure 1.

As expected, the distribution of wealth of the 1960–64 birth cohort is to the left of the population distribution for young ages and it moves to the right as the cohort ages. Nevertheless, it is clear that changes in within-cohort wealth ranks always correspond to meaningful changes in wealth levels, more so for later periods.

FIGURE B.1. Population vs within-cohort ranks for the 1960–64 birth cohort

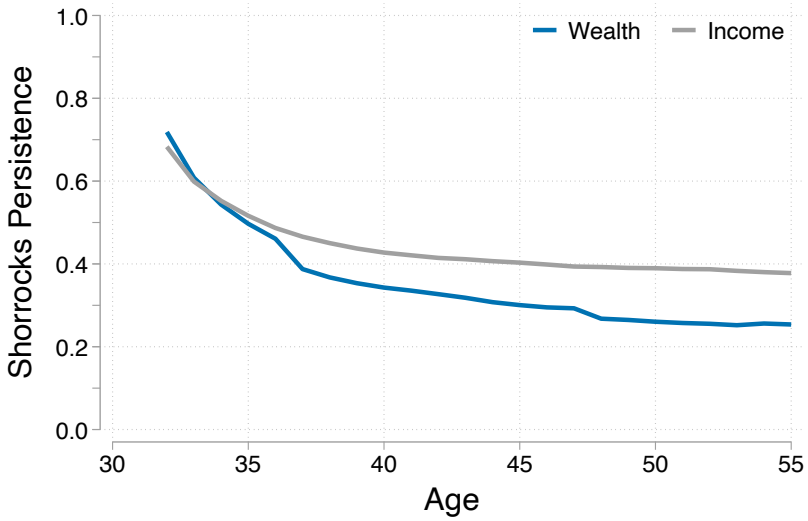


Notes: The figure shows the average wealth rank in the Norwegian population of individuals in the 1960–64 birth cohort at different ages. A value in the 45° line implies that the within-cohort coincides with the population rank. Numbers are average wealth holdings in 2019 US dollars by percentile.

B.2. Alternative measure of rank persistence

Figure B.2 reports the Shorrocks' (1978) persistence measure as an alternative measure of intragenerational rank persistence to rank autocorrelations, as used in Figure 2.

FIGURE B.2. Intragenerational persistence of wealth ranks

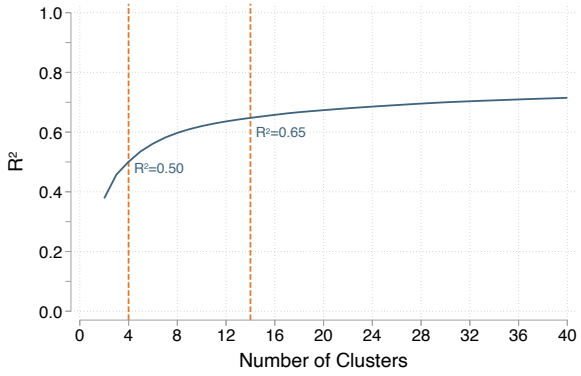


Notes: The figure plots the Shorrocks' (1978) persistence measures for intra-generational mobility of wealth and income, corresponding to the share of individuals who at time t are in the same cohort wealth quintile of wealth or income, respectively, that they were in 1993.

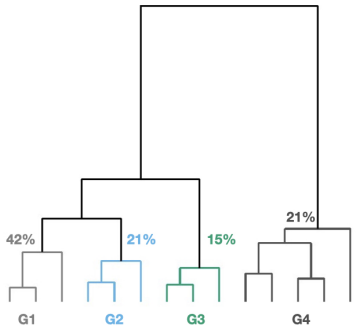
Here, the Shorrocks' (1978) measure is defined as the share of individuals who at time t are in the same wealth quintile of wealth or income, respectively, that they were in 1993. Figure B.2 confirms the rank autocorrelation results in Figure 2. Life-cycle wealth mobility increases along the life cycle, with the share of individuals in the same cohort wealth quintile as in their early thirties dropping to around 30 percent by age 40 before stabilizing at approximately 25 percent from their late forties onward. We also find that there is less life-cycle income mobility than in wealth using Shorrocks' (1978) measure.

FIGURE C.1. Choice of number of groups

(A) Share of total variation explained



(B) Dendrogram



Notes: Panel A shows the share of the variation explained as the number of cluster increases. Panel B presents the dendrogram of the hierarchical agglomerative clustering procedure as executed on the balanced sample for the 1960–64 birth cohort. The dendrogram shows the tree of clusters up to a hierarchy of $G = 14$ groups. The tree shows how groups are merged as the clustering procedure recursively reduces the number of groups.

Appendix C. Additional results on typical trajectories

C.1. R^2 measures for partitions

We operationalize our choice of the number of groups, G, using the R^2 measure, corresponding to the share of the variation in individual trajectories explained by the cluster average (or typical) trajectory. For a partition $\mathcal{G}_G = \{g_i\}_{i=1}^N$ over the set of individuals with $G > 1$ groups, this measure is

$$R^2 = 1 - \frac{\sum_{i,t}\left(y_{i,t} - \bar{y}_t^{g(i)}\right)^2}{\sum_{i,t}\left(y_{i,t} - \bar{y}\right)^2}, \tag{C.1}$$

where $y_{i,t}$ is the wealth rank of individual i at time t, $\bar{y}_t^{g(i)}$ is the average rank for the individual i's group at that same time, and $\bar{y}$ is the average rank of individuals across the (balanced) sample.

Figure C.1A presents the R^2 for the partitions produced by our hierarchical clustering algorithm for $G = 1, \ldots, 40$. With four groups, we capture 50 percent of the variation in wealth ranks trajectories, while keeping the exercise parsimonious. Going from $G = 4$ groups to $G = 14$ groups (the thinnest level of granularity shown in Figure C.1B) only increases the R^2 from 50 percent to 65 percent.

The nested nature of our clustering algorithm allows us to transparently illustrate how our baseline choice of $G = 4$ groups affects our findings. The hierarchy of groups is summarized by the dendrogram in Figure C.1B. Each of the small branches at the bottom represents smaller clusters obtained at the step $G = 14$, and the tree shows how they are recursively aggregated up (agglomerated) by the procedure into a single cluster ($G = 1$). We highlight in different colors the four baseline groups that we select. We can then directly assess how sensitive these groups are to alternative values of G. For instance, $G = 5$ splits group 4 in two, while $G = 3$ would merge groups 1 and 2. It is therefore straightforward to see which groups are closest through the lens of the procedure. We study the wealth trajectories of the main subgroups of our baseline groups in Section 9.2 and Appendix C.4.

Between and within R^2. The R^2 defined in (C.1) is a function of the variation between groups, captured by the *between* R^2, and within groups, captured by the *within* R^2.

$$R^2_{between} = 1 - \frac{\sum_i (\tilde{y}_i - \bar{\tilde{y}}^{g(i)})^2}{\sum_i (\tilde{y}_i - \bar{\tilde{y}}_i)^2} \qquad \text{between } R^2 \qquad \text{(C.2)}$$

$$R^2_{within} = 1 - \frac{\sum_i \sum_t ((y_{i,t} - \tilde{y}_i) - (\bar{y}^{g(i)}_t - \bar{\tilde{y}}^{g(i)}))^2}{\sum_i \sum_t ((y_{i,t} - \tilde{y}_i) - (\bar{y} - \bar{\tilde{y}}_i))^2} \qquad \text{within } R^2 \qquad \text{(C.3)}$$

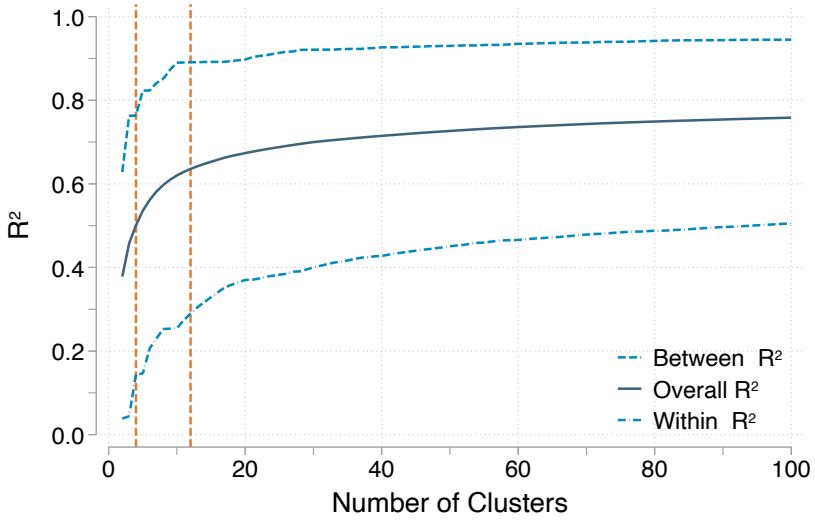
The between R^2 captures how dissimilar the typical (average) trajectories are across groups, and it is a good indicator of whether the groups are meaningfully different. It measures the share of the cross-sectional variation in ranks (having averaged over the longitudinal dimension of the panel) explained by the clusters' typical trajectories. For this we define $\tilde{y}_i = \sum_t y_{i,t}/T$ as the within person average rank, $\bar{\tilde{y}}_i$ its average across individuals, and $\bar{\tilde{y}}^{g(i)}_i$ its average for cluster $g(i)$. The between R^2 for $G = 4$ groups is close to 80 percent.

The within R^2 captures how much underlying heterogeneity is there inside the groups. It measures the share of the variation in ranks along the longitudinal dimension of the panel explained by the clusters' typical trajectories. For this we define the deviation of an individual's rank in time t relative to the own average rank as $y_{i,t} - \tilde{y}_i$ and contrast it with the population wide average deviation in ranks $\bar{y} - \bar{\tilde{y}}_i$. This gives a measure of the total *within* variation. The cluster's deviation, $\bar{y}^{g(i)}_t - \bar{\tilde{y}}^{g(i)}$, where $\bar{y}^{g(i)}_t$ is the cross-sectional average of ranks for cluster $g(i)$ in time t and $\bar{\tilde{y}}^{g(i)}$ is its average over time is used to construct a measure of the explained within variation produced by the clusters' typical trajectories. The within R^2 for $G = 4$ groups is close to 15 percent, but it more than doubles for $G = 14$ groups.

Figure C.2 presents the three R^2 measures for the largest 100 clusters produced by our agglomerative hierarchical clustering algorithm. With 100 clusters it is possible (but not optimal) to group individuals based on their initial (or final) wealth rank and trace their trajectories. Most of the increase in explanatory power takes place with the

first 20 clusters. As expected, the vast amount of variation in the data is hard to capture, reflected by the lower value of the within R^2 even when clusters increase. In contrast to this, the between R^2 reaches 0.8 with 4 clusters and close to 0.9 with 14, reflecting the fact that even with few clusters, the wealth trajectories they generate are significantly different and capture a large share of the overall variation in wealth mobility.

FIGURE C.2. R^2 Measures, up to 100 clusters



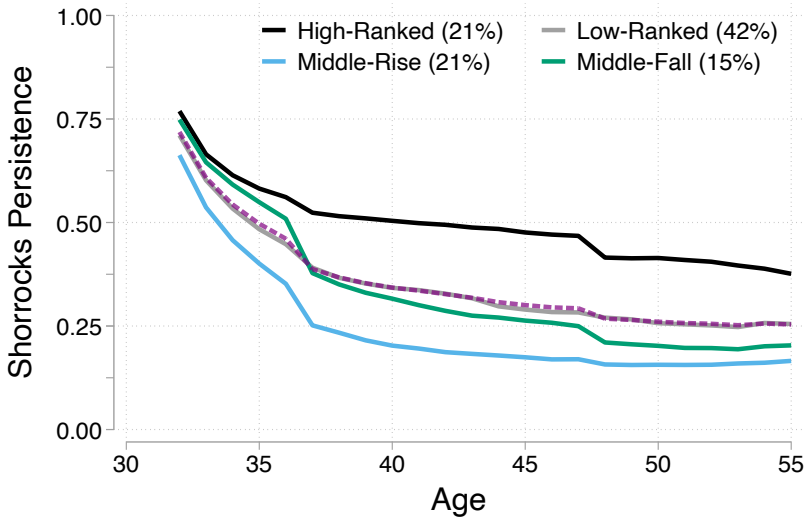
Notes: The figure plots the share in the variation of wealth rank trajectories explained by the partitions induced by agglomerative hierarchical clustering algorithm for $G = 1, \ldots, 100$ groups. The overall R^2, defined in equation (C.1), and also presented in Figure C.1A. The between R^2, defined in equation (C.2), captures the share of variation across clusters. The within R^2, defined in equation (C.3), captures the average share of variation within each clusters.

C.2. Intragenerational mobility by group: Shorrocks

Analogously to Figure 5, we can also decompose the Shorrocks mobility measure reported in Figure C.3 by group. We compute the Shorrocks index of mobility given by the share of individuals in each group who remain in their initial quintile of the cohort's wealth distribution as they age. Consequently, the population level measure is a weighted average of our group-specific measures (where the weights correspond to group size).

The intragenerational mobility trends of all the groups follow the same qualitative pattern as the decomposition of rank persistence we focus on in the main text (Figure 5). Each group exhibits distinct patterns of wealth mobility. The high-ranked group displays higher levels of intragenerational persistence (lower mobility) than the population average. They are almost twice as likely to remain in the same quintile they were in at age 30 by age 55. By contrast, mobility is higher for the groups of risers and fallers,

FIGURE C.3. Intragenerational persistence across group: Shorrocks index



Notes: The Figure plots the intragenerational Shorrocks index for our four main groups. The pooled cohort-level persistence measure is shown in dashed lines.

reflecting large changes in ranks over the life cycle, with risers experiencing the highest mobility.

C.3. Further group characteristics

We provide here additional results characterizing the four main groups we identify in Section 5.

Figure C.4 complements Figure 6 by providing life-cycle profiles for the asset and debt levels of the four groups. The main takeaway is that the differences in net-worth reported in Figure 4 mostly reflect differences in assets across groups. The magnitude of debt and dispersion of debt are much lower than those of assets.

Since housing wealth is the main asset in the portfolio of all groups (Figure 6A), we also separately analyze the accumulation of non-housing wealth, defined as wealth excluding the value of individuals' primary residence. Figure C.5 depicts the evolution of non-housing wealth in the four main groups both in ranks (Panel A) and in levels (Panel B). Differences in non-housing wealth are subdued relative to differences in total wealth. Strikingly, there is no gap in non-housing wealth between the Risers and Fallers before age 40, and the gap we document in total wealth can therefore be ascribed almost entirely to differences in housing wealth earlier in life, as shown in Figure 6B. Besides, Figure C.5B also shows that non-housing wealth is mostly accumulated by the group of individuals permanently at the top of the wealth distribution and the group rising through the wealth distribution. Individuals in the other two groups barely accumulate any non-housing wealth over their lifetime.

Figures C.6 and C.7 complement Figure 7 by providing the share of individuals in each group that receive benefits from major public programs and the average life cycle trajectories of employment and self-employment income for the four groups. With respect to public programs, low-earning, low-wealth individuals are more likely to be beneficiaries. For example, 10 percent of the individuals in the low-ranked group receive unemployment benefits at age 40, and more than 20 percent receive disability benefits, compared with less than 5 and 10 percent, respectively, in the high-ranked group. There is also a clear trend of a reduction in unemployment benefits and an increase in disability benefits as the cohort ages.

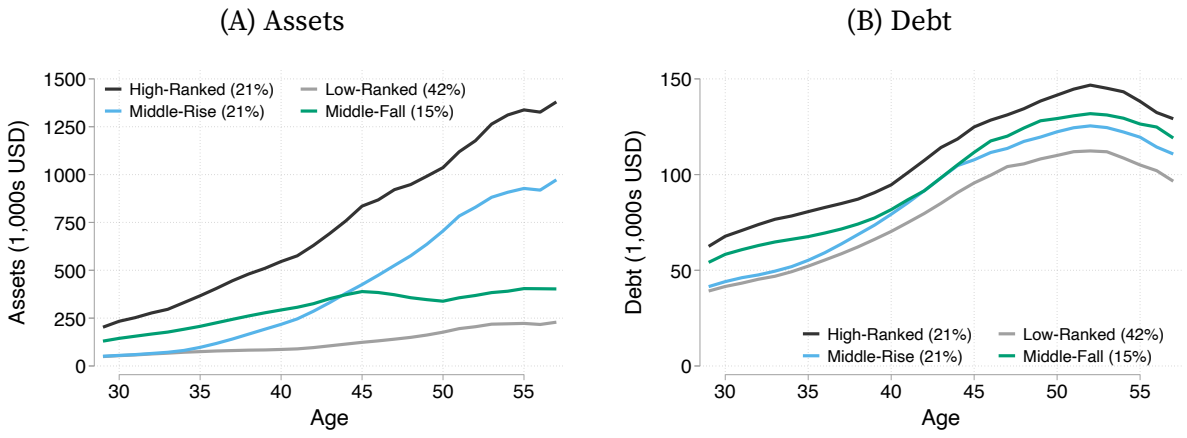
Figure C.7A shows that the differences in employment income among groups follow the same trends as the differences in overall income shown in Figure 7B but are not as large. The discrepancy lies in the differences in self-employment and capital income. Figure C.7B shows the life-cycle profile of self-employment income. It exhibits the same pattern of having the high-ranked at the top and the risers overtaking the fallers as they age. These graphs report unconditional averages within each group, and therefore reflect extensive margin differences in employment and self-employment (see Figure 7D). The remaining income differences between groups come from the life-cycle profile of capital income. Capital income is low for all groups except for the high-ranked group. The high-ranked group has, on average, 10,000 to 15,000 US dollars of capital income (with fluctuations across years), while all other groups have at most 5,000 dollars (this is the case of the risers after age 50).

Figure C.8 presents the life cycle profile of household wealth ranks for the individuals in each of the four main groups of the 1960–64 birth cohort. Household wealth ranks have the same qualitative (and even quantitative) behavior as individual wealth ranks. To construct households, we match individuals in the 1960–64 birth cohort with their spouses or partners, corresponding to their civil status of married or cohabitating. The matching uses the complete population file. Individuals are then assigned their households' wealth before we compute within-cohort household wealth ranks.

The role of human wealth for mobility. The differences in income trajectories between mobility groups point toward a relevant role for *human wealth*, that is, the discounted value of future realized labor market income, in driving the trajectories of financial wealth we have documented so far. Accordingly, we look at the evolution of the sum of financial and human wealth for information on how individuals save out of the labor market income they effectively receive. We present the results in Figure C.9.

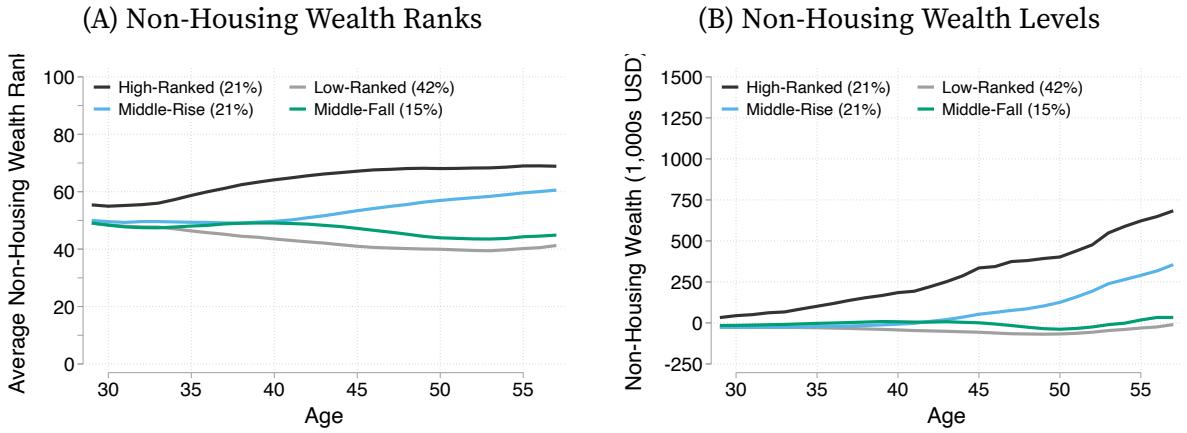
Taking into account human wealth closes the gap in initial financial wealth between risers and fallers. That is, the risers' higher income trajectories fully compensate for their lower initial financial wealth. Interestingly, the sum of financial and human wealth remains relatively constant until age 45 when their trajectories diverge, reflecting different savings patterns (to which we return in Section 7). The diverging pattern of these wealth trajectories coincides with the crossing pattern in financial wealth and the increase of non-property assets in the risers' portfolio documented above. Human wealth plays a lesser role for the mobility of high- and low-ranked groups because of

FIGURE C.4. Asset and debt trajectories by group



Notes: Panels A and B plot, respectively, average asset and debt trajectories in 2019 US dollars of the four main groups presented in Figure 3.

FIGURE C.5. Wealth Beyond Housing



Notes: The figures present characteristics of the non-housing wealth of the four main groups presented in Figure 3. Panel A reports the evolution of non-housing wealth ranks, while panel B reports the average levels of non-housing wealth. Non-housing wealth is defined as total wealth minus the value of primary residences.

their respective high- and low-income trajectories. Nevertheless, the distance between groups narrows when taking into account human wealth, reflecting differences in income trajectories not accounted for by differences in financial wealth.

TABLE C.1. Composition of wealth and income groups

(a) Composition of wealth groups in terms of income groups

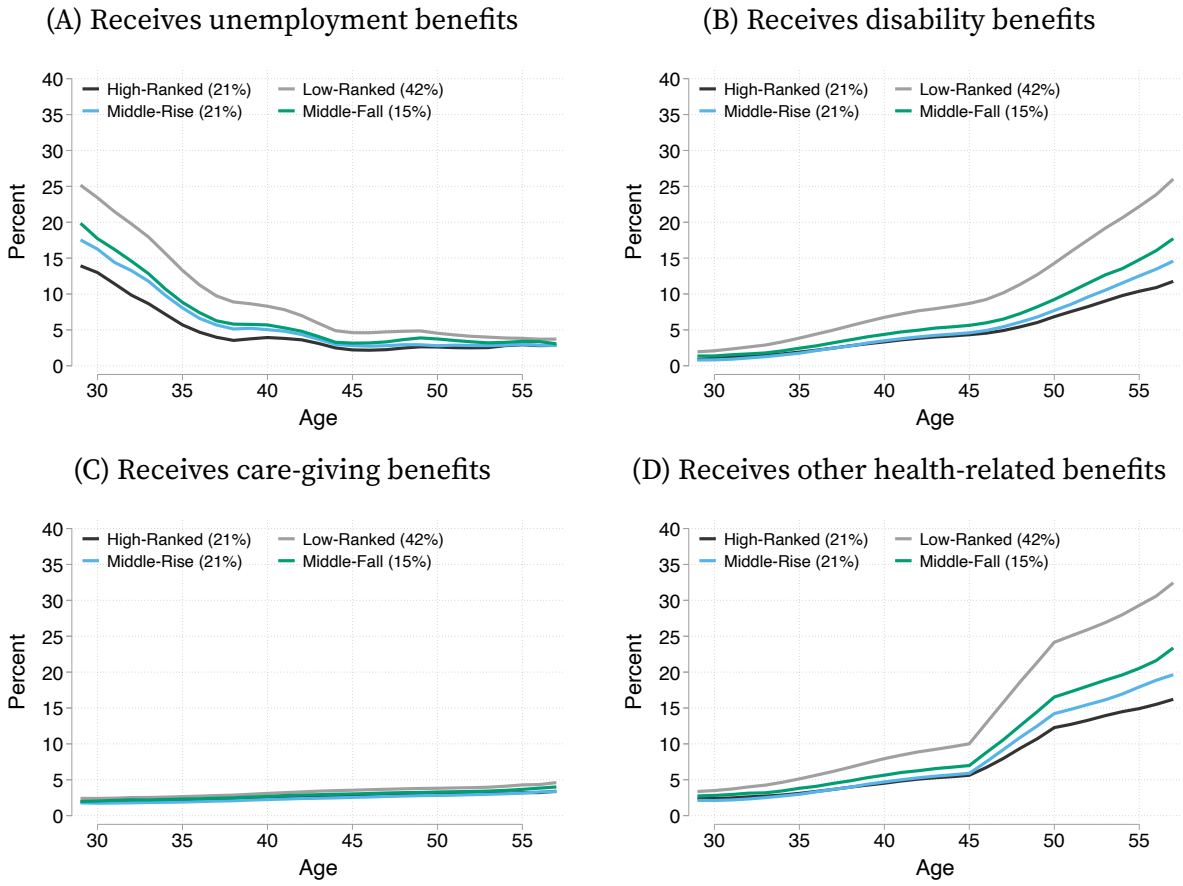
	Inc. Gr. 1	Inc. Gr. 2	Inc. Gr. 3	Inc. Gr. 4	Pop. Share
High-Ranked	17.5%	19.1%	27.1%	36.4%	21.0%
Low-Ranked	43.8%	23.3%	20.8%	12.1%	42.7%
Middle-Rise	27.3%	23.1%	26.4%	23.2%	21.0%
Middle-Fall	29.4%	23.6%	27.5%	19.6%	15.4%

(b) Composition of income groups in terms of wealth groups

	Inc. Gr. 1	Inc. Gr. 2	Inc. Gr. 3	Inc. Gr. 4
High-Ranked	11.3%	17.9%	23.4%	37.0%
Low-Ranked	57.2%	44.2%	36.3%	25.0%
Middle-Rise	17.6%	21.6%	22.8%	23.5%
Middle-Fall	13.9%	16.2%	17.4%	14.5%
Pop. Share	32.6%	22.4%	24.3%	20.7%

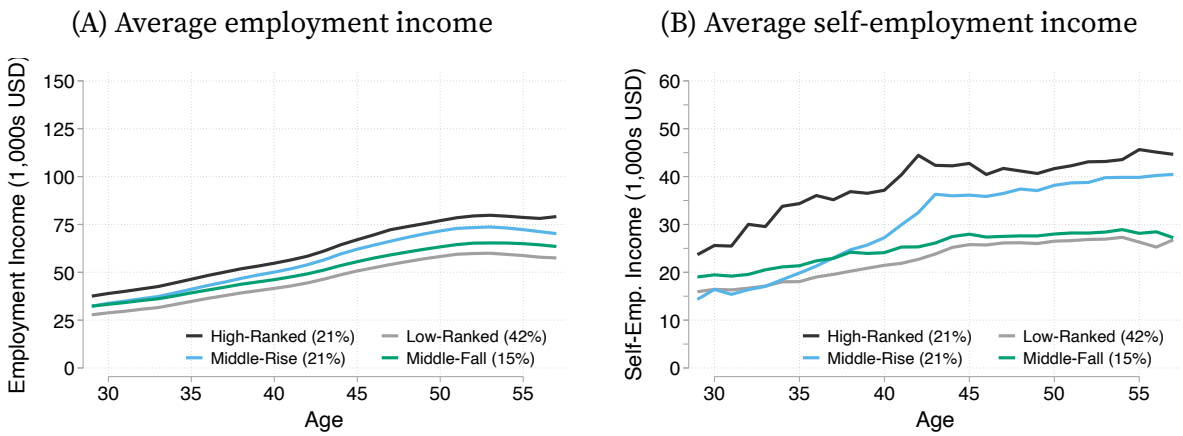
Notes: Panel a reports the composition of each of the wealth groups in Figure 3 in terms of the income groups in Figure 9. Panel b reports the composition of each of the income groups in Figure 9 in terms of the wealth groups in Figure 3.

FIGURE C.6. Benefit receipt by group



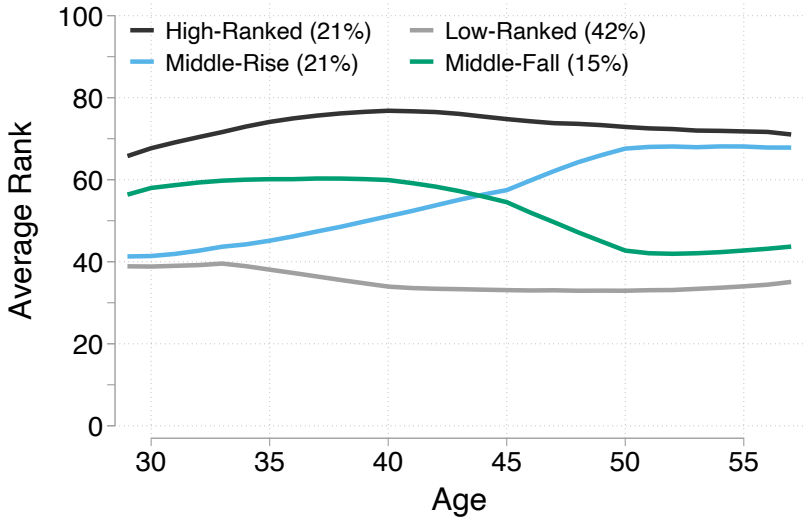
Notes: The figures plot the share of individuals by age that receive benefits from public programs for the four main groups presented in Figure 3. We consider four types of benefits: unemployment; disability; transfers related to care-giving including those for parents of young children, elderly care, and sick people; and health related benefits encompassing transfers to individuals undergoing treatment or those experiencing major health issues themselves or for close relatives. All numbers are in percentage points.

FIGURE C.7. Employment and self-employment income by group



Notes: Panels A and B plot, respectively, the average employment and self-employment income trajectories in 2019 US dollars for the individuals in each of the four main groups presented in Figure 3. The average is taken over all the individuals in the group and therefore is a result of the intensive and extensive margin of employment and self-employment.

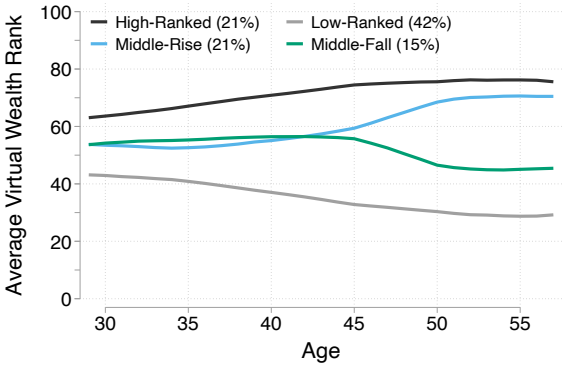
FIGURE C.8. Household wealth ranks by group



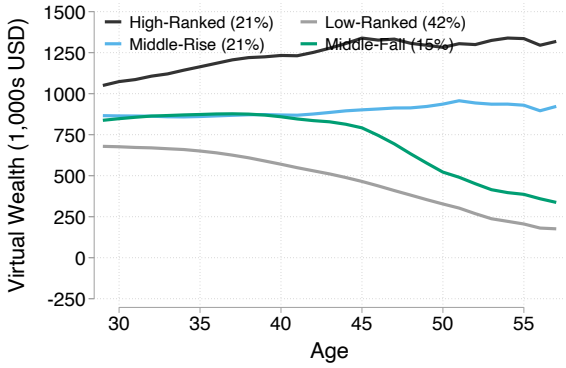
Notes: The figure plots the rank of household wealth for the individuals in each of the four main groups presented in Figure 3. To construct households, we match individuals in the 1960–64 birth cohort with their spouses or partners, corresponding to their civil status of married or cohabitating. The matching uses the complete population file. Individuals are then assigned their households' wealth before we compute within-cohort household wealth ranks.

FIGURE C.9. The Role of Human Wealth for Mobility

(A) Wealth + Human Wealth Ranks



(B) Wealth + Human Wealth Levels



Notes: The figures present ranks (Panel A) and levels (Panel B) of the sum of financial and human wealth of the four main groups presented in Figure 3. Human wealth is defined as the discounted value of future labor income calculated from the realized income trajectories of individuals obtained from the tax registry. We discount future income using the average return on net worth for Norway, 3.21 percent, reported in (Fagereng et al. 2020, Table 3).

C.4. Group characteristics for subgroups

We provide here additional results characterizing the subgroups of our four main groups, described in Section 9.2. We describe the main patterns below.

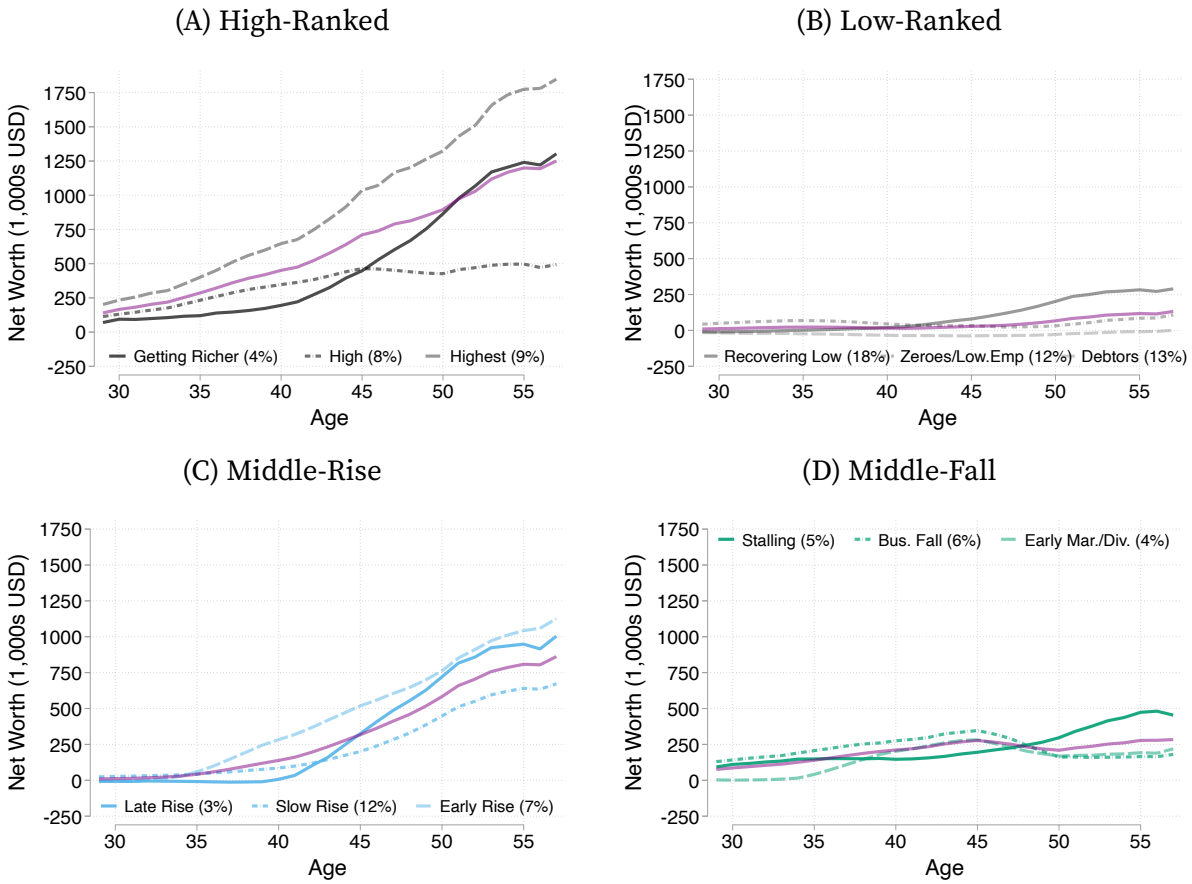
The typical mobility patterns of these groups are presented in Figure 15 in the main text. The corresponding typical wealth trajectories are in Figure C.10. The common scale in the figure highlights the vast differences in net-worth across subgroups. These are much larger than the differences across groups because of the "highest" group among the high-ranked. The largest differences in debt are consequently concentrated within the high-ranked group, with the largest increase in wealth being presented by the "getting richer" subgroup. Among the risers there is convergence between the late and early risers despite starting to accumulate wealth at different ages. For the low-ranked, the lowest subgroup is the only one that has consistently negative net-worth, hence the moniker of "debtors". The "recovering" subgroup differs from the others in its accumulation of property (as we see in Figure C.12). The differences among the fallers are less pronounced. The early marriage/divorce group stands out because of the accumulation of property (again, Figure C.12). This group's moniker follows from its household formation and dissolution dynamics.

Figures C.11 and C.13 present the composition of subgroup's portfolio and income, respectively. Once again the largest differences are in the high-ranked group, where the "highest" subgroup has a much larger share of private assets and stocks than any other group. Correspondingly, it also has the highest share of capital and self-employment income among all subgroups. Interestingly, the "getting richer" subgroup is closer to the subgroups of the risers in terms of its portfolio and income composition, even though it does have larger shares of private assets and stocks and the income share of capital and self-employment income increases faster for this group than for the risers. Most other groups have high shares of property assets and labor income throughout.

In terms of the prevalence of self-employment, Figure C.14 shows that it is concentrated among the subgroups that start highest in the wealth distribution: the "highest" and "high" subgroups of the high-ranked and the "business fall" subgroup of the fallers. These differences align with the patterns of portfolio and income composition discussed above.

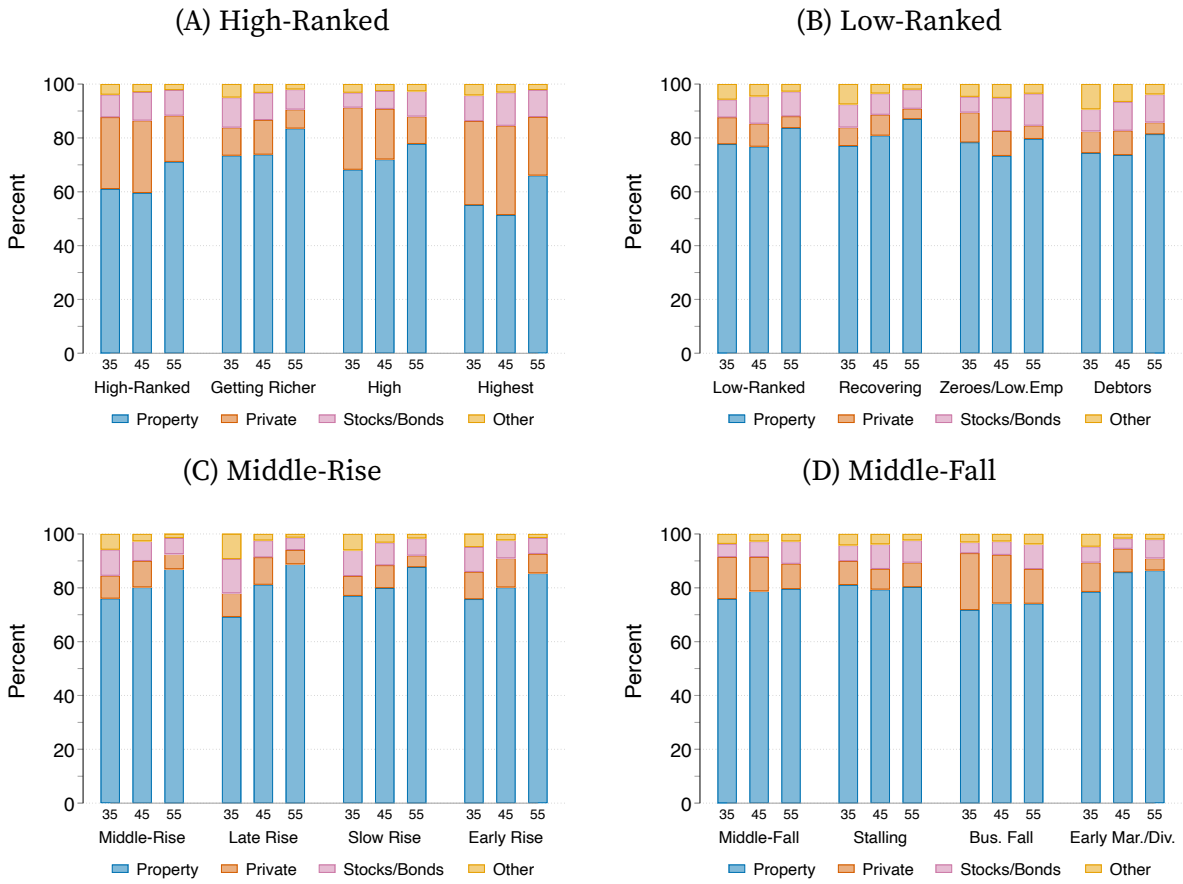
Finally, Figure C.15 reports educational attainment for all four main groups and their subgroups. Risers (and their subgroups) have the highest educational attainment, even when compared with the high-ranked subgroups. Tellingly, the subgroup with the next highest education is the "getting richer" subgroup of the high-ranked; this is the subgroup that most relies on labor income among those at the top. The contrast between these groups and the subgroups of the fallers is also telling. None of the subgroups of the fallers have postgraduate rates above 7 percent, while the risers have rates of at least 10.5 percent. The differences are even larger for undergraduate degrees.

FIGURE C.10. Wealth levels by subgroup



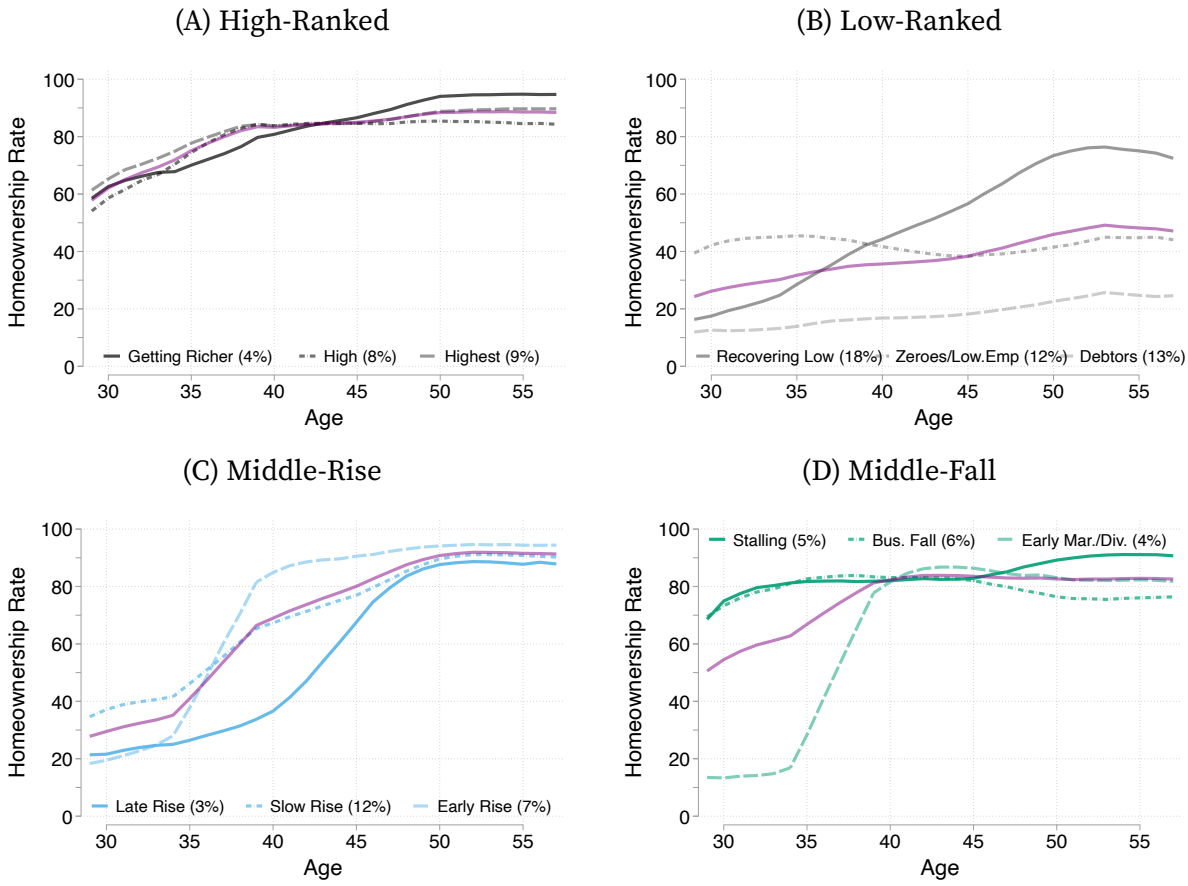
Notes: The figure plots the average wealth level in thousands of US dollars in each clustered subgroup against the cohort's average age. All individuals belong to the 1960–64 birth cohort. For each subgroup we additionally report the major group's values. This corresponds to Figure 4 in the main text.

FIGURE C.11. Portfolio composition by subgroup



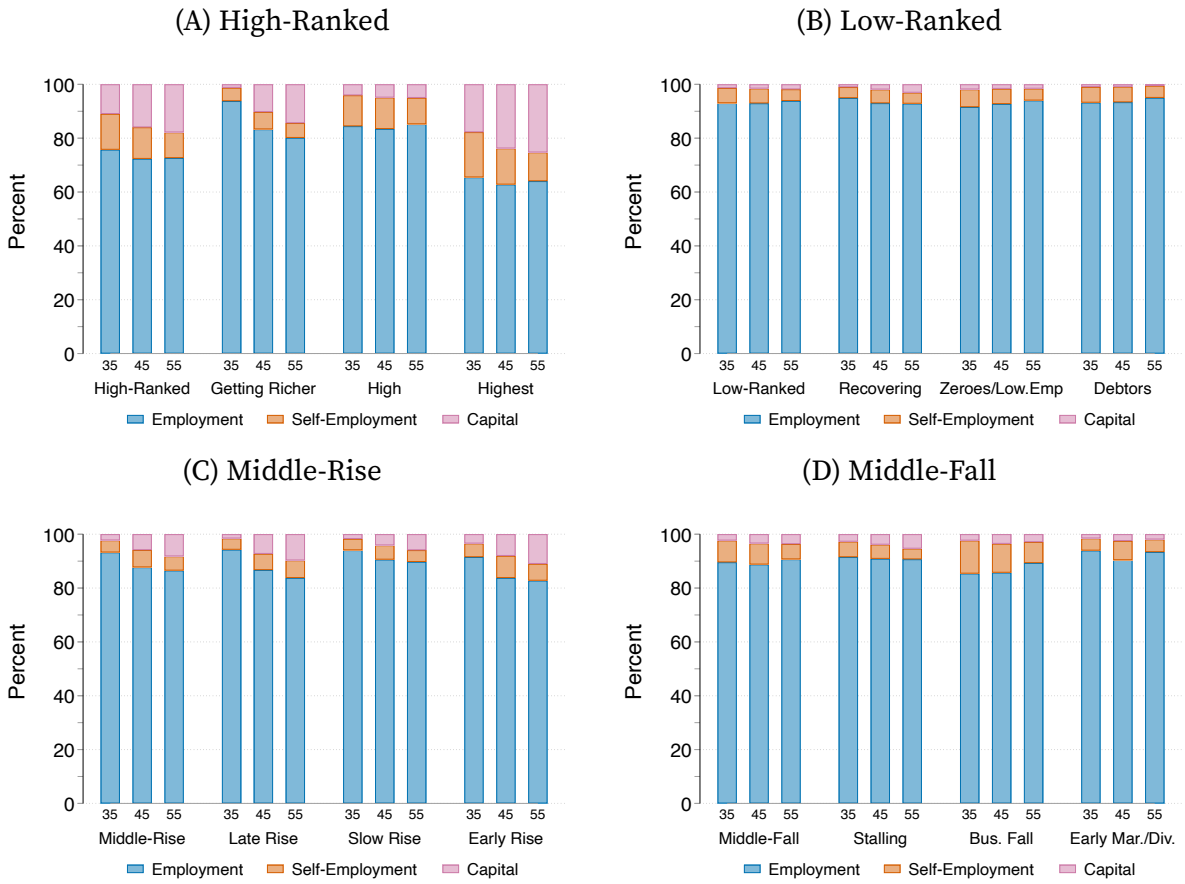
Notes: This figure reports the share of assets accounted for by property, privately held assets, financial assets, and other assets, defined as the total value of each asset class divided by the total assets within a group. For each subgroup, we additionally report the major group's values. This corresponds to Figure 6A in the main text.

FIGURE C.12. Homeownership rates by subgroup



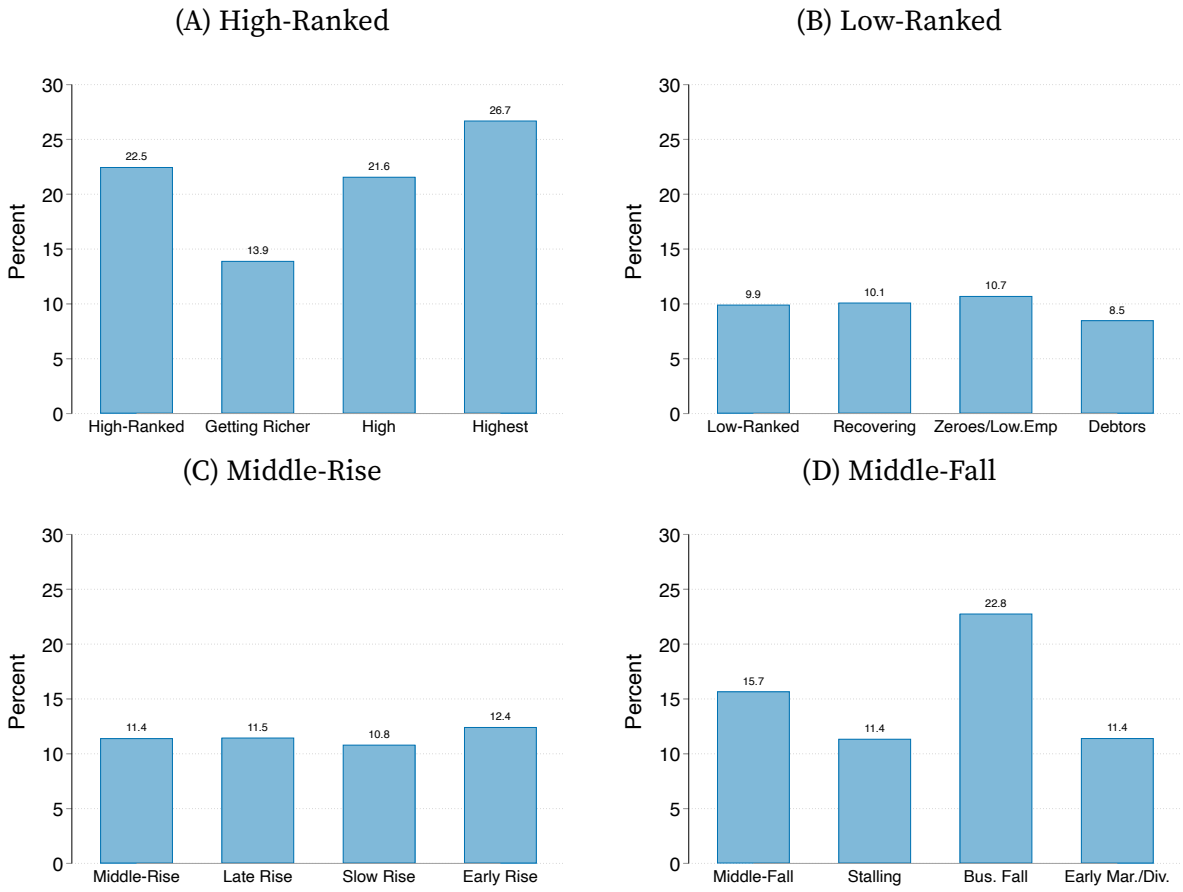
Notes: The figure plots the share of individuals who are homeowner in each clustered subgroup against the cohort's average age. For each subgroup we additionally report the major group's values. This corresponds to Figure 6B.

FIGURE C.13. Income composition by subgroup



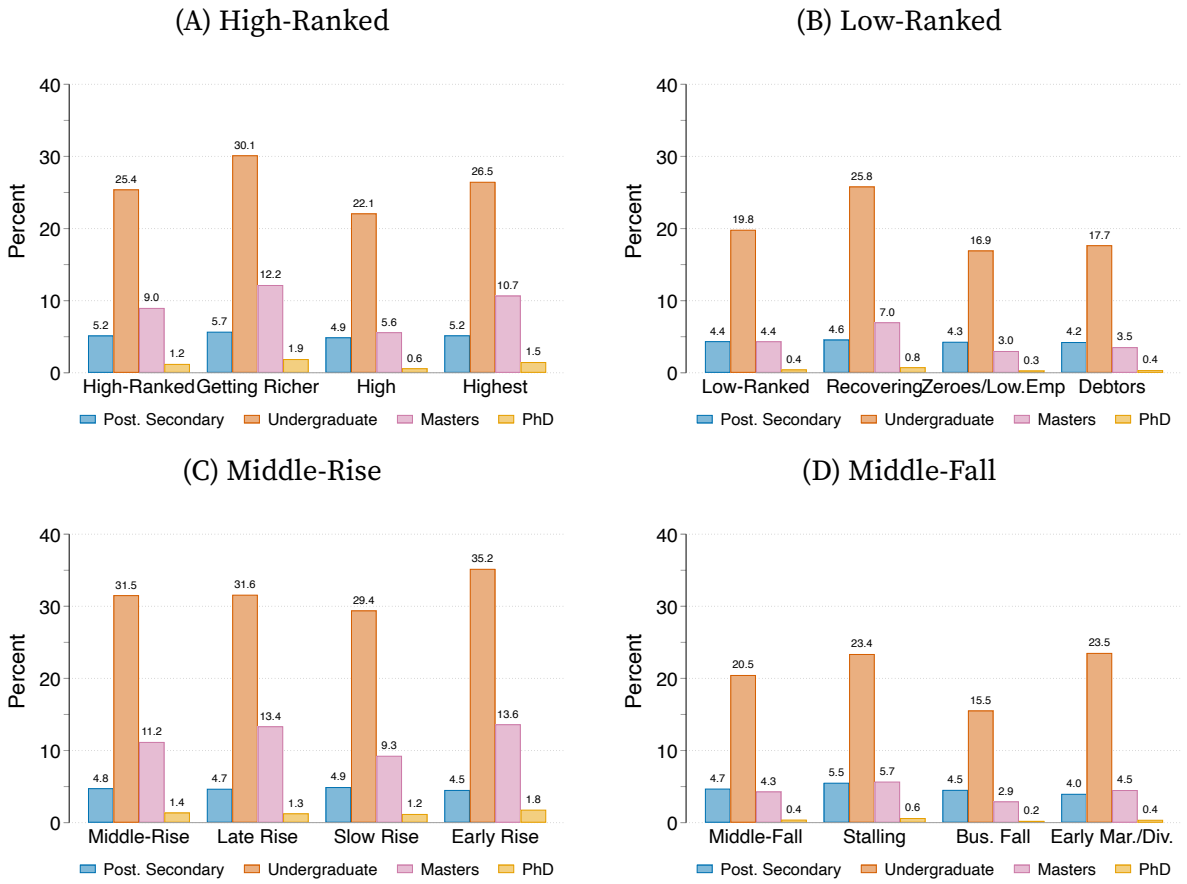
Notes: This plots the share of each group's income accounted for by employee, self-employment, and capital income. For each subgroup, we additionally report the major group's values. This corresponds to Figure 7C.

FIGURE C.14. Self-employment rates by subgroup



Notes: This figure plots the share of individuals in each group with self-employment income. For each subgroup, we additionally report the major group's values. This corresponds to Figure 7D.

FIGURE C.15. Educational attainment by subgroup



Notes: This plots the share of individuals in each group achieving different levels of educational attainment For each subgroup, we additionally report the major group's values. We discuss the categorization of educational attainment in Section 8.

Appendix D. Proofs for Section 7

Assumptions on income process. We assume that skill heterogeneity, I, belongs to a finite number of values $I \in \{\underline{I}, \ldots, \bar{I}\}$[63] and that the function mapping idiosyncratic shocks into age-varying efficiency units is log-linear in the idiosyncratic shock, $\ln y_t(z) = \alpha_t + \beta_t \ln z$, where the coefficients capture age-dependent non-linearity.

These two technical assumptions are satisfied for every version of the model where permanent skill heterogeneity does not enter the age-profile of efficiency units (see, for example, Guvenen, Karahan, Ozkan, and Song 2021; Guvenen, Kambourov, Kuruscu, Ocampo, and Chen 2023). We continue to work with the model without taxation of incomes to simplify the exposition, which is without loss when taxes are approximated by a power function (Bénabou 2002; Heathcote, Storesletten, and Violante 2014). This preserves the log-linearity of after-tax incomes in permanent skill heterogeneity and their shocks. We discuss implications of relaxing this assumption in Appendix D.2.

Parallel log-income profiles. We define total income, $\text{inc}_{i,t}$, for each agent i at time t as the product of efficiency units and permanent skill-heterogeneity so that

$$\ln \text{inc}_{i,t} \;=\; \ln I_i + \alpha_t + \beta_t \ln z_{i,t}. \tag{D.1}$$

Log-linearity follows immediately as a consequence of the assumptions on the income function. Relying on this log-linear expression gives the following expressions for the mean of age t income conditional on skill heterogeneity

$$E_t[\ln \text{inc}_{i,t} \,|\, I] \;=\; \ln I_i + \alpha_t + b_t E_t[\ln z_{i,t} \,|\, I] \;=\; \ln I_i + \alpha_t, \tag{D.2}$$

where we assume $E_t[\ln z_{i,t}] = 0$ without loss of generality. In this equation the age-efficiency profile α_t determines the overall shape of average income and skill heterogeneity determines the location. Consequently, for any two values of skill heterogeneity we have

$$E_t[\ln \text{inc}_{i,t} \,|\, I] - E_t[\ln \text{inc}_{i,t} \,|\, I'] \;=\; \ln I - \ln I' \tag{D.3}$$

which is independent of t, implying parallel income profiles.

D.1. Proof of Proposition 1

Parallel log-wealth and consumption profiles. It is easy to verify that the model we present satisfies Assumptions (1)-(3) in Straub (2019) (who considers a more general

[63] This is relevant for the structure of our proof, however, it is true by construction in all numerical implementations of the model we consider where the solution necessitates a (potentially very fine) discrete grid.

framework).[64] The fact that log-wealth age profiles are parallel and shifted by permanent income I follows directly from Lemma 1 in Straub (2019).

Formally, Lemma 1 in Straub (2019) directly implies $F(a \mid t, I) = F(a \times \frac{I'}{I} \mid t, I')$, where $F(a \mid t, I)$ denotes the distribution of wealth (a) at age t for permanent income group I. We can further normalize $I = 1$, which is without loss of generality since the age-income profile (D.1) can always be shifted by a constant, and therefore $F(a \mid t, 1) = F(a \times I \mid t, I)$.

Average wealth in permanent income group I is then

$$\begin{aligned} E[a \mid t, I] &= \int_0^{\bar{a}(t,I)} a \; f(a \mid t, I) da \\ &= \int_0^{\bar{a}(t,I)/I} [\tilde{a} \times I] \; [f(\tilde{a} \times I \mid t, I) \times I] d\tilde{a} \\ &= I \times \int_0^{\bar{a}(t,1)} \tilde{a} \; f(\tilde{a} \mid t, 1) d\tilde{a} \\ &= I \times E[a \mid t, 1], \end{aligned} \tag{D.4}$$

where the second line uses the change of variable $a = \tilde{a} \times I$ and the third line uses $f(a \mid t, 1) = f(a \times I \mid t, I)I$. The notation $\bar{a}(t, I)$ denotes the upper bound of the wealth distribution for age t and permanent income group I.

Let $p_I(t) = \lim_{a \downarrow 0} F(a|t, I)$ denote the mass point associated with the binding borrowing constraint. It follows that this mass point is the same size for all permanent income groups, $p_I(t) = p_{I'}(t) = p(t)$, so that the borrowing constraint binds for the same fraction of individuals in each permanent income group.

It follows that for all $t > 0$ and $a > 0$,

$$F(a \mid t, 1, a > 0) = \frac{F(a \mid t, 1)}{1 - p(t)} = \frac{F(a\frac{I'}{1} \mid t, I')}{1 - p(t)} \implies H(\ln a \mid t, 1) = H(\ln a + \ln I' \mid t, I'), \tag{D.5}$$

where $H(\cdot)$ denotes the distribution of log wealth and h its corresponding density.

Average log-wealth in permanent income group I can then be written

$$\begin{aligned} E[\ln a \mid t, I] &= \int \ln a \; h(\ln a|t, I) \; d \ln a \\ &= \int \ln a' \; h(\ln a' + \ln I|t, I) \; d \ln a' \; + \; \ln I \\ &= \int \ln a' \; h(\ln a'|t, 1) \; d \ln a' \; + \; \ln I \\ &= E[\ln a'|t, 1] \; + \; \ln I, \end{aligned} \tag{D.6}$$

where the second line uses the change of variable $\ln a = \ln a' + \ln I$ and the third line

[64] Like Straub (2019) we also require an additional uniqueness assumption that rules out degenerate income risk.

uses the equality in D.5. It follows directly that for any two values of skill heterogeneity we have

$$E[\ln a \mid t, I] - E[\ln a \mid t, I'] = \ln I - \ln I', \tag{D.7}$$

which is independent of t, implying parallel wealth profiles.

The same steps apply for consumption. This shows the second consequence we state.

Consistency of hierarchical clustering. Consider panel data on income and wealth trajectories for the population with T observations in the time dimension and N individuals. Assume the income observed in the data, $inc_{i,t}$, is the product of efficiency units and permanent skill heterogeneity. Thus, it corresponds to the model object above. Likewise, assume the panel data on wealth is produced by the policy function of the model. There are $G > 0$ values of permanent skill ($I = 1, \ldots, G$).

Proposition 1 shows that, if the data is partitioned by permanent skill heterogeneity, it generates parallel profiles. We now prove that clustering will partition the data by permanent skill heterogeneity. As above, we can exploit the log-linearity of income to express the variance of age t income conditional on skill heterogeneity as

$$V_t[\ln \text{inc}_{i,t} \mid I] \;=\; b_t^2 V_t[\ln z_{i,t} \mid I] \;=\; b_t^2 \sigma_z^2, \tag{D.8}$$

and

$$V_t[\ln a_{i,t} \mid I] \;=\; V(\ln I_i \mid I) + V_t[\ln a_t(\frac{a}{I}, 1, z) \mid I] \;=\; \sigma^2(t). \tag{D.9}$$

We denote the collection of these averages (for either income or wealth) as $\boldsymbol{\mu}_I$ which is a T length vector of average log outcomes conditional on a value of permanent skill-heterogeneity. The T-by-T matrix $\boldsymbol{\Sigma}_I$ will similarly denote the variance-covariance matrix.

Without loss of generality, we order the groups by the variability of their features

$$\text{tr}\left(\boldsymbol{\Sigma}_j\right) \leq \text{tr}\left(\boldsymbol{\Sigma}_{j'}\right) \quad \text{for} \quad j < j' \quad (\leq G), \tag{D.10}$$

and, thus, G denotes the group with the largest variability and group 1 denotes the group with the smallest variability.

The asymptotic consistency of the classifier then follows from applying Theorem 5 in Egashira, Yata, and Aoshima (2024), that establishes the following necessary and sufficient condition for a classifier using Ward's distance metric:[65]

$$\underset{T\to\infty}{\text{limsup}} \left(\frac{|\ \text{tr}\left(\boldsymbol{\Sigma}_1\right) - \text{tr}\left(\boldsymbol{\Sigma}_G\right)\ |}{\min_{j<j'} ||\ \boldsymbol{\mu}_j\ - \boldsymbol{\mu}_{,'}\ ||^2} \right) \;<\; \min_{j<j'} \left\{ N_j N_{j'} / \left(N_j + N_{j'} \right) \right\}. \tag{D.11}$$

[65] For alternative metrics the right hand side expression is replaced by the value of one half.

Intuitively, this states that the gap in average outcomes must be larger than the additional variability induced by observing an increasing number of features (in our case a longer trajectory of income or wealth). Moreover, in order to discriminate consistently this must be larger the smaller the smallest group is (notice that only the time dimension is varying in this condition, so that the result holds for any finite and fixed number of individuals, see also footnote 24).

In the class of homothetic buffer-stock savings models presented here, this condition is satisfied immediately for either clustering on observed income or wealth when skill-heterogeneity is not degenerate (see footnote 64). The numerator on the left-hand-side is always zero as there is no difference in variability across groups and the denominator is strictly positive for any value of T.

With this condition being satisfied, the agglomerative hierarchical clustering algorithm classifier is asymptotically consistent for groups defined by their skill-heterogeneity. This concludes the proof of Proposition 1.

COROLLARY A1 (**Ordered expected ranks**). *The income and wealth ranks of each permanent income group are also ordered, that is,*

$$E_t[rank_{i,t} \,|\, I^{'}] < E_t[rank_{i,t} \,|\, I] \;\forall t, I^{'} < I, \tag{D.12}$$

where $rank_{i,t}$ corresponds to the income or wealth rank of individual i at age t.

PROOF. Without loss of generality, we set $I^{'} = 1$ in what follows. We show here that the inequality above holds for wealth ranks. The result for income ranks is immediate because income is homogeneous of degree 1 in permanent income I and, thus, so is the *distribution* of income.

Let $Q_t(a) = F^{-1}(a|t)$ denote the population quantile function, so that $\text{rank}_{i,t} = Q_t(a_i)$. This is a monotonic function in a for all values of t. There is no conditioning on permanent income group. Additionally, Let X and Y be random variables that correspond to the wealth of individuals in permanent income groups $I^{'}$ and I, respectively. It is from the ordering of these distributions that we derive the order for the expected rank across groups. The distribution function of X is $F_X(a) = F(a|t, I^{'})$ and the distribution of Y is $F_Y(a) = F(a|t, I) = F(a|t, 1)$. It follows from Lemma 1 in Straub (2019) (scaling of distributions) and from the fact that $I > 1$ that

$$F_X(a) = F_Y\left(\frac{a}{I}\right) \longrightarrow F_X(a) < F_Y(a) \;\forall a, \tag{D.13}$$

therefore X (with higher permanent income) dominates Y in the usual stochastic order.

Expected group ranks then satisfy

$$E_t[rank_{i,t} \,|\, I^{'}] = E_t[Q_t(X)] < E_t[Q_t(Y)] = E_t[rank_{i,t} \,|\, I], \tag{D.14}$$

where we use the fact that X and Y have the distributions defined above. This inequality follows immediately from the definition of stochastic dominance and the fact that Q is

an increasing function. □

Thus, there is also no mobility in ranks in the sense that expected ranks are ordered across groups, with the order reflecting their permanent income. Two groups who differ in their permanent income can never switch expected ranks, even though their rank trajectories are not necessarily parallel.

This additional result has a clear implication for the behavior of wealth ranks that we can verify with our data by comparing Figures 3 and 9B. The two figures have markedly different mobility patterns. This indicates that the wealth mobility patterns in the population do not correspond to those of permanent income, in violation of homothetic buffer-stock savings, even when the expected wealth ranks across permanent income groups do not cross (as expected).

D.2. Relaxing homotheticity and implications for interpretation

In the proof of Proposition 1 we assume homotheticity of total income in skill heterogeneity and homotheticity of preferences, so that the policy functions are also homothetic. However, these assumptions may not hold if, for example, skill-heterogeneity affects the life-cycle profile of efficiency units or the income risk faced by groups (Gourinchas and Parker 2002). We now discuss the implications of relaxing these assumptions.

We first express total income as

$$\ln \text{inc}_{i,t} = \ln I_i + \alpha_t(I) + \beta_t(I) \ln z_{i,t}, \tag{D.15}$$

which allows for skill-heterogeneity to affect the life-cycle profile of efficiency units directly as well as the income risk faced by groups. It follows that

$$E_t[\ln \text{inc}_{i,t} \mid I] = \ln I_i + \alpha_t(I) + b_t(I) E_t[\ln z_{i,t} \mid I] = \ln I_i + \alpha_t(I), \tag{D.16}$$

and the variance can be expressed as[66]

$$V_t[\ln \text{inc}_{i,t} \mid I] = b_t(I)^2 V_t[\ln z_{i,t} \mid I] = b_t(I)^2 \sigma_z^2(I). \tag{D.17}$$

In this case, we do not necessarily have parallel groups and $\alpha_t(I)$ can be such that two groups intersect and cross (as could be the case for low- and high-education groups, with the former starting from a higher income level to be surpassed by the latter later in life). However, we will have approximately parallel groupings if the difference in permanent income is larger than the difference in the transitory income components across groups, that is, $\ln I - \ln I'$ is large relative to $\alpha_t(I) - \alpha_t(I')$.

Verifying, the conditions for an asymptotically consistent classification in equation (D.11) are also no longer immediate. This is because the variance of data features differ

[66] Given we allow for the transmission of persistent income to differ, we do not normalize the variance of z to be constant across groups

across groups, and the additional variability can confound the classifier's ability to correctly discriminate based on small differences in the means. Intuitively, we require that differences in variability across groups be dominated by the differences in the average location of their income profile to retain the consistency result presented above. Therefore, we require the group differences are meaningful in a sense equation (D.11) makes formal.[67] This requirement of meaningful group differences is shared by work applying similar algorithms in the context of income dynamics (see Section 9.4).

We can apply the same logic to study the implications of clustering on wealth profiles. To do so, we re-express the log profile in terms of mean-zero deviations from homotheticity:

$$\ln a_{i,t} \;=\; = \;\ln I_i + \overline{\mathcal{A}}_t(I) + \varepsilon_{i,t}, \tag{D.18}$$

where the distribution of ε captures both persistent income shocks and differences across groups in saving decisions. When deviations are small, $\overline{\mathcal{A}}_t(I) \approx \overline{\mathcal{A}}_t$, log-wealth profiles are close to parallel. Similarly, when group differences satisfy equation (D.11), e.g. if $\sigma_\varepsilon(t, I) \approx \sigma_\varepsilon(t)$, then clustering on wealth is produces groups that are partitioned by permanent differences in skill-heterogeneity.

If, however, deviations from homotheticity are large, for example, by producing large heterogeneity in savings rates across the wealth distribution or because rates-of-return are heterogeneous, then clustering will not produce groups partitioned by skill-heterogeneity. *Far from being a negative result, this is precisely the logic we appeal to in the main text.* We view the fact that clustering on wealth does not produce groups defined by their permanent skill-heterogeneity as evidence that the data is not generated by behavior that is close to the homothetic model, lending support to presence of meaningful and persistent differences in rates of return and savings rates that drive wealth accumulation over the life cycle.

[67] For example, if $b_t(I) = b_t$ and differences in the variance of shocks across groups are small, $\sigma_z^2(I) \approx \sigma_z^2$, then it is easy to see that the numerator of equation (D.11) is small relative to the denominator.

Appendix E. Additional results on ex ante analysis

E.1. The role of parental education, sex, and birthplace

Field of study. Panel A of figure E.1 shows the average partial effects associated with the field of education for individuals with at least an undergraduate degree.

Parental education. Panel B of figure E.1 shows the average partial effects associated with the highest level of parental educational attainment. We find that the effects are muted throughout, with risers being the only group with a differential educational attainment among parents. The parents of risers are more highly educated and thus risers are more likely to have parents with postgraduate or PhD degrees.

Sex. Panel C of figure E.1 shows the average partial effects associated with sex at birth. We find that men are substantially more likely to be in the high-ranked group and less likely to be risers by approximately 10 percentage points. We also find that they are slightly more likely to be fallers and less likely to be in the low-ranked group by approximately 5 percentage points.

Birthplace. Panel D of Figure E.1 reports the average partial effect estimates for the place of birth indicators. We find a positive effect of being born in Oslo or another larger Norwegian city on the probability of being in the high-ranked and middle-rise wealth mobility groups. Although significant, these effects are smaller in magnitude (about 5 percentage points) than those we find for parental wealth and education.

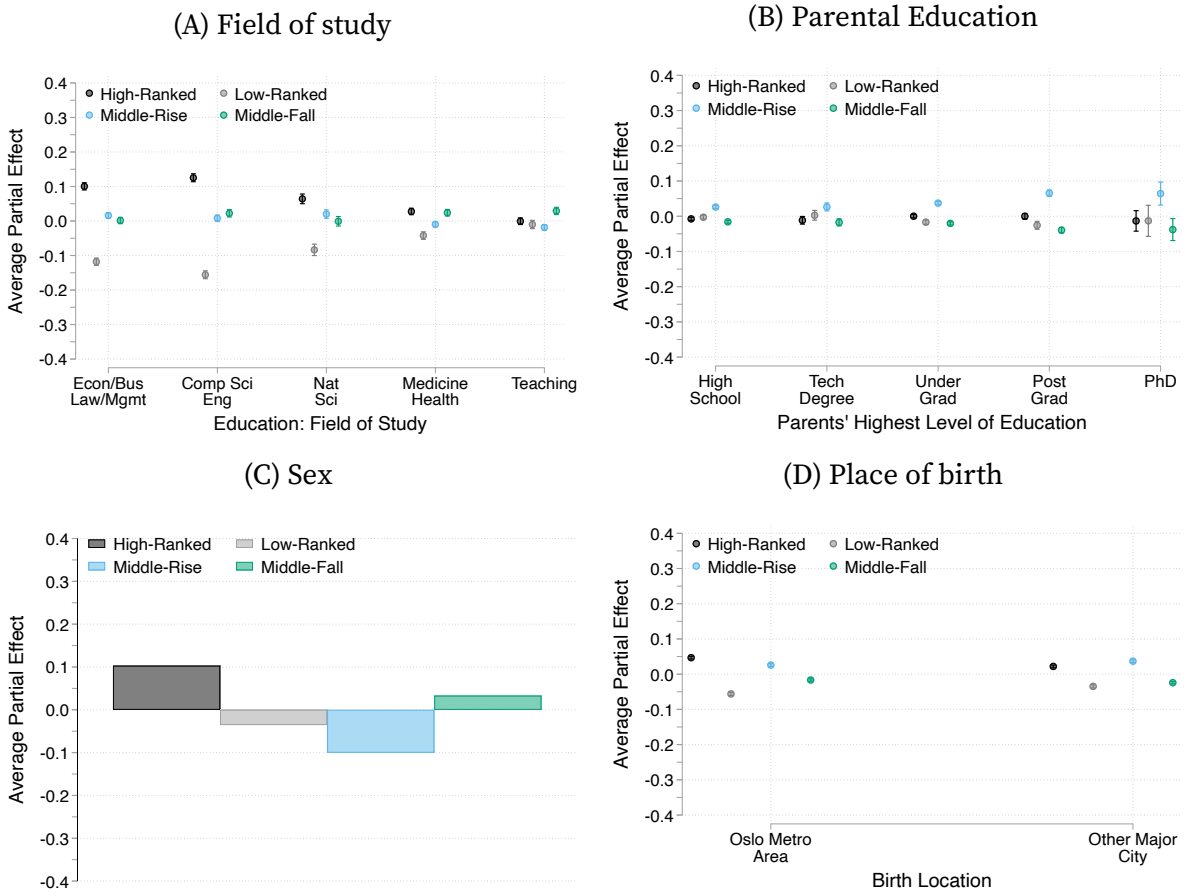
E.2. Additional covariates: Own and parental background

In order to investigate the robustness of the role of ex ante determinants we consider three alternative specifications. In the first, we include additional proxies for the dynamics of parental wealth, namely, whether parents owned a business. In the second, we include additional information on the wealth and portfolios of individuals in 1993. We include a series of *own wealth* ventile fixed effects, as well as binary indicators for whether an individual was a homeowner or owned a business in 1993. Finally, we include both groups of variables in a third specification.

Overall, we find little predictive power for parental businesses or whether an individual owned a home or business in 1993, conditional on the values of parental wealth, education, sex, and birthplace. Instead, we find a large role for own wealth. We interpret the important role for an individual's own wealth in 1993 as consistent with the patterns of segmented mobility we document in Section 5. We report the average partial effects for these additional controls in Figure E.2. The results correspond to the third specification with all the covariates.

When we include these additional controls, the explanatory power of our classifier increases almost four-fold and is driven almost entirely by initial wealth in 1993 (see

FIGURE E.1. Demographics and the probability of group assignment

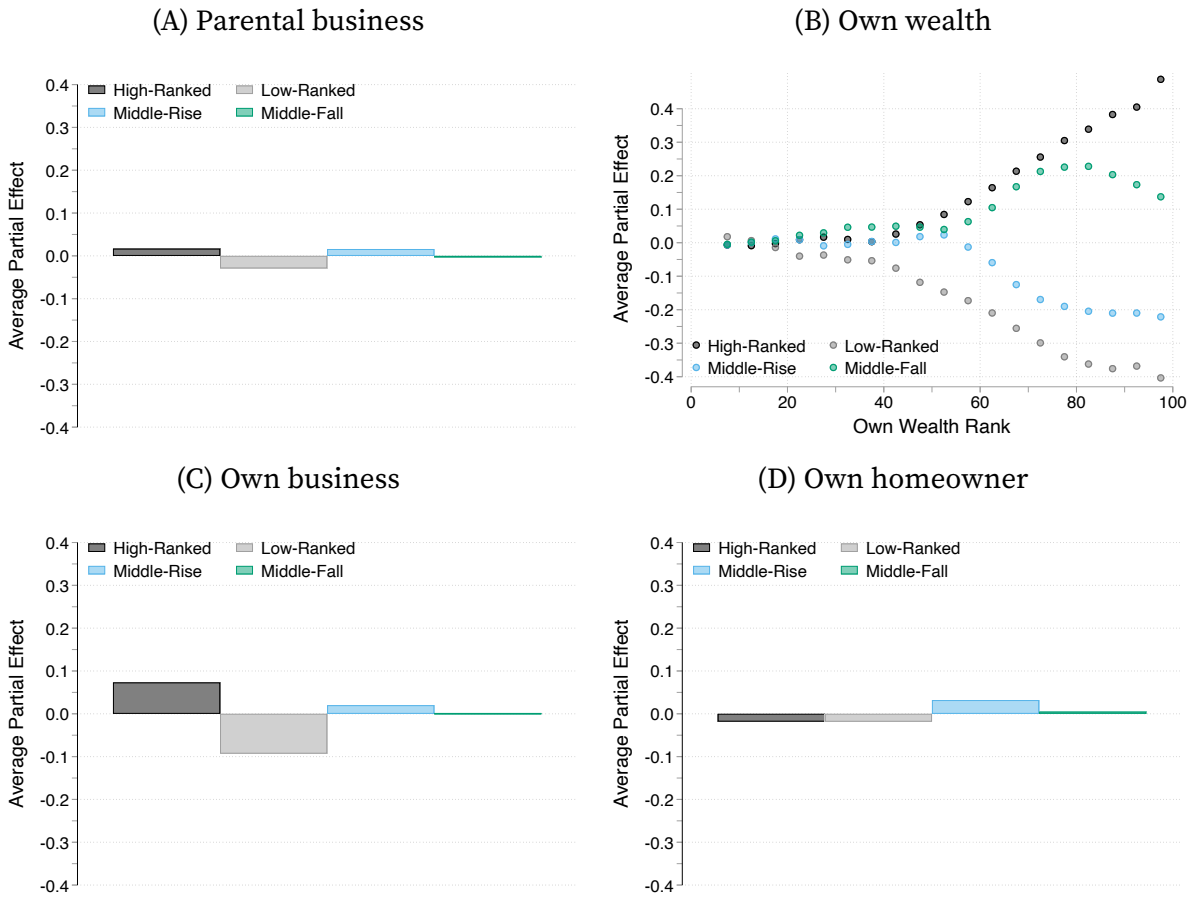


Notes: Panel A plots the average partial effect of field of study (for those with technical degree or above) relative to a humanities degree. Panel B plots the average partial effect of parental educational attainment relative to compulsory schooling age. Panel C plots the average partial effect of men relative to women. Panel D plots the average partial effect of urban areas relative to rural areas. We construct the average partial effect by integrating over the empirical joint distribution of other covariates. We report point estimates separately for each outcome, the probability of being assigned to each of our four groups, along with their 95 percent confidence intervals.

Table E.2). Once we know where an individual begins, we are able to accurately categorize the typical trajectory because the existence of segmented mobility implies that initial wealth is an accurate discriminator of outcomes over the whole life cycle.

As we describe in Section 8, once we include this additional information on an individual's initial wealth, the role of parental wealth and education declines. In general, the point estimates for the average partial effects decline by 25–40 percent. Figure E.3 reports the estimated average partial effects for our two key drivers (parental wealth and educational attainment) under this alternative specification. Although initial wealth absorbs some of the explanatory power of these variables, the predictions

FIGURE E.2. Parental portfolio, own wealth and the probability of group assignment



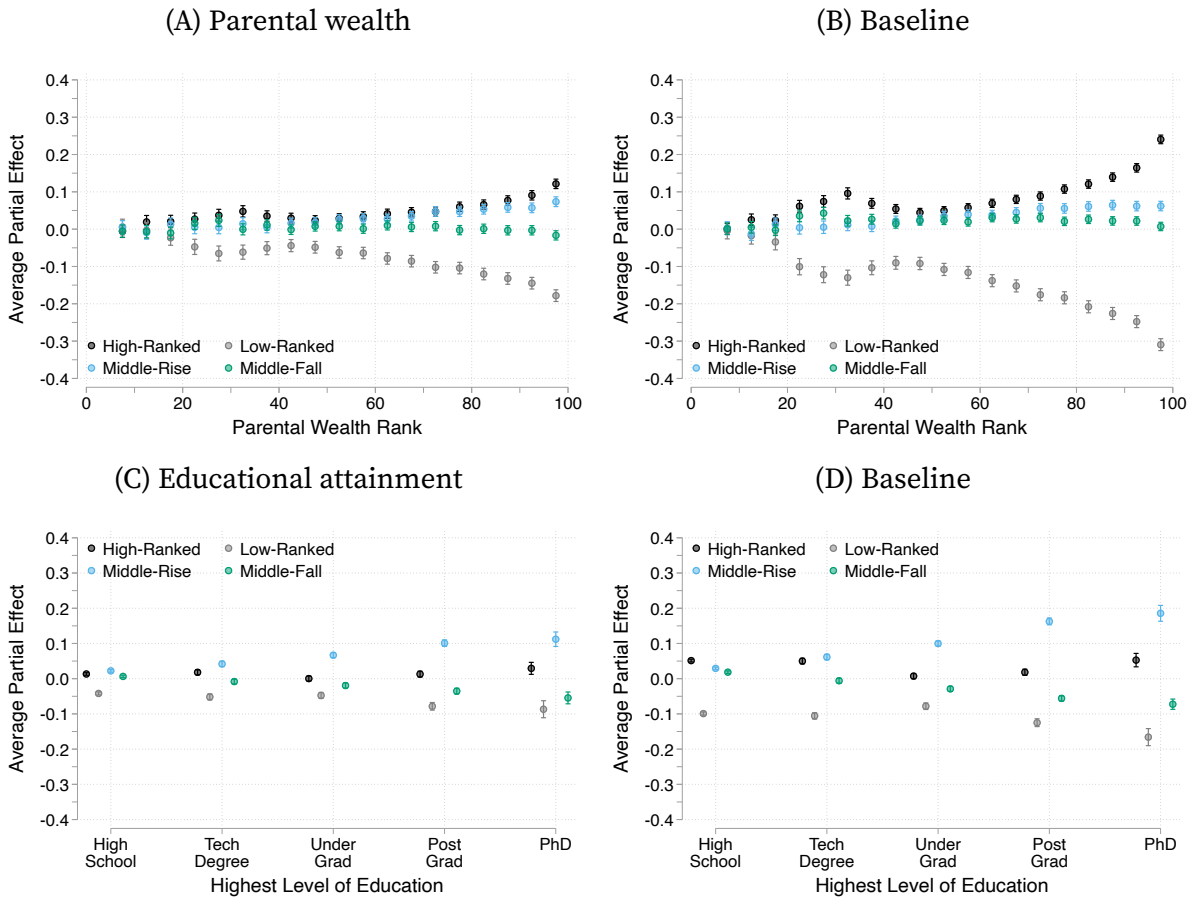
Notes: The figures plot the average partial effect on the probability of belonging to each of the four main groups of whether parents have a business in 1993 (panel A), the individual's own wealth ventile in 1993 (panel B), whether the individual own a business in 1993 (panel C), and whether the individual owns a house in 1993 (panel D). The effects of own wealth ventiles are reported relative to being in the bottom ventile of the distribution in 1993. All the results correspond to the specification of the multinomial logit in (6) with the addition of the variables described in the text. We construct the average partial effects by integrating over the empirical joint distribution of all other covariates. We report point estimates separately for each outcome, the probability of being assigned to each of our four groups.

remain significant and display the same qualitative patterns—confirming the relative importance of education and parental wealth among risers and the high-ranked.

E.3. Relative predictive power of ex ante characteristics

Our results show that parental background and the individual's initial conditions play an important role in determining group membership. We now explore how each set of ex ante covariates in Equation (6) helps to explain the variation across groups.

FIGURE E.3. Robustness of parental wealth rank, educational attainment and the probability of group assignment



Notes: The figures contrast the average partial effects on the probability of belonging to each of the four main groups of parental wealth and education with and without additional covariates for parental business, own wealth in 1993, and business and homeownership in 1993. Panels A and B present the effects of parental wealth ventiles in 1993, relative to being born to parents in the bottom ventile of the distribution. Panels C and D present the effects of educational attainment, relative to compulsory schooling age. We construct the average partial effect by integrating over the empirical joint distribution of other covariates. We report point estimates separately for each outcome, the probability of being assigned to each of our four groups, along with their 95 percent confidence intervals.

We use two measures to gauge the predictive power of the ex ante characteristics of individuals. First, we measure the share of variation explained using the Distance-Weighted Classification Rate

$$1 - \frac{\sum_{i=1}^{N}\sum_{k=1}^{G}\widehat{\Pr}(g=k|X_i)d(g(i),k)}{\sum_{i=1}^{N}\sum_{k=1}^{G}\widehat{\Pr}(g=k)d(g(i),k)}, \tag{E.1}$$

where $d(g, g')$ corresponds to Ward's distance metric in equation (3).

The distance-weighted classification rate, which is bounded between 0 and 1, corresponds to the average implied distance between an individual's true group and their predicted group $\widehat{\Pr}(g = k|X_i)$ weighted against a naive predictor $\widehat{\Pr}(g = k)$ that uses a homogeneous random assignment. As the distances between disjoint groups are positive, the numerator of the fraction in equation (E.1) can be interpreted similarly to the residual sum of squares in the coefficient of determination, while the denominator can be interpreted as the total sum of squares. Consequently, a value of one implies perfect classification, while a value of zero implies that the covariates contain no information. Because this measure considers the distance across groups, it penalizes more strongly classifying a low-ranked as a high-ranked rather than as a faller.

The second measure we use is the Unweighted Classification Rate

$$\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{G}\widehat{Pr}\left(g = k \,|\, X_i\right)\ \mathbb{1}[g(i) = k]. \tag{E.2}$$

As with the first measure, this measure is between 0 and 1 by construction. Unlike the first measure, this measure only cares about the rate of correctly classified individuals relative to a naive predictor $\widehat{\Pr}(g = k)$ that uses a homogeneous random assignment and does not depend on the type of misclassification that takes place. In exchange, its units are immediately interpretable as the (extra) share of correctly classified individuals with respect to random assignment.

Tables E.1 reports the total contribution from our four groups of exante regressors to the distance-weighted and unweighted classification rates. We also report in each table a decomposition of the partial contribution of each regressor using a Shapley-Owen decomposition (Shorrocks 2013). This decomposition allows us to calculate a single value per covariate category that is permutation-invariant and additively-decomposable despite the nonlinearity of the object being decomposed (in this case, the classification rates). We describe the Shapley-Owen-Shorrocks decomposition in Appendix F.

Consistent with the average partial effects we report in Figures 12A and **??**, we find that parental background and education account for the majority of the model's explanatory power, with a much more limited role for gender and place of birth. These results hold for the weighted and unweighted classification rates. We interpret this result as evidence that observable ex ante characteristics matter for individual wealth mobility over the life cycle.

The relative contribution of parental wealth to correctly classifying individuals is also high—accounting for over 40 percent of the fit of the model in classifying individuals across groups (Appendix E.3 Table E.1). This contribution takes into account how parental wealth varies jointly with other individual characteristics. The contribution varies when classifying individuals of different groups. It is 45 and 46 percent for the high- and low-ranked groups, respectively, while it is only 20 percent for risers and 4.3 percent for fallers.

Education accounts for almost 40 percent of the ability of the model for classifying

TABLE E.1. Predictive power of ex ante characteristics

(a) Share of distance variation explained by variable (pp)

Group	All Effects	Partial Contribution			
		Parent	Education	Sex & Birthplace	Par. Edu.
All	5.91	41.28	39.11	13.47	6.14
High-Ranked	7.91	44.99	29.33	20.51	5.16
Low-Ranked	7.00	46.01	41.57	4.85	7.57
Middle-Rise	4.63	20.08	54.94	20.82	4.16
Middle-Fall	0.28	4.30	15.50	82.18	-1.98

(b) Share of individuals correctly classified (pp)

Group	Random	All Effects	Partial Contribution			
			Parent	Education	Sex & Birthplace	Par. Edu.
All	29.33	3.15	34.31	40.41	19.23	6.04
High-Ranked	21.03	4.40	41.08	23.57	30.52	4.83
Low-Ranked	42.51	3.34	45.28	44.33	3.24	7.15
Middle-Rise	20.91	3.52	12.58	50.14	33.59	3.70
Middle-Fall	15.55	0.73	11.88	47.63	25.83	14.67

Notes: The tables report the distance-weighted and unweighted probability of belonging to an individual's true group, across the cohort in row "All" and conditional on being in each of the four main groups in the remaining rows. The distance corresponds to the measure in equation E.1. The combined explanatory power of all covariates is reported in column "All". The remaining columns report the partial contribution of each variable category as a share (in percentage points) of their combined explanatory power in column "All." The classification model corresponds to our estimated multinomial logit presented in equation (6). The explanatory power is computed relative to a naive random classification, the classification rates for this model are reported in the bottom table under the "Random" column. The partial contribution of each variable category is obtained through the Shapley-Owen decomposition, which averages across permutations of decompositions and sums to the total contribution, as described in Appendix F.

individuals across groups, so that together with parental wealth they make up for most of the explanatory power of initial characteristics (Appendix E.3, Table E.1). Education variables are most relevant for the explanatory power of the model in classifying risers, where it accounts for almost 55 percent of correct classification rate; it accounts for almost 30 percent for the high-ranked group.

Sex and birthplace explain less than 20 percent of the fit of the model in classifying individuals, mostly because of sex differences in the groups' composition.

Relative to the results we report in the main text, Table E.1 reveals an important additional pattern. Although, on average, the discriminating power of education is lower than that of the parental background, its ability to classify individuals is much more consistently spread across groups. By contrast, parental background is most effective at correctly classifying those at the extremes of the distribution (the high- and low-

ranked); it only has limited informational content for predicting those who will rise or fall through the churn in the middle of the distribution. We view this as highlighting an important notion of equality of opportunity: extreme comparisons point to *inequality* of opportunity, but there is more *equality* of opportunity in the middle of the distribution.

We find that, on average, these covariates explain around 6 percent of our distance measure. The share of variation explained by the variables in (6) is similar in magnitude to the R^2 values reported in intergenerational estimates of the rank correlation in wealth—specifically to those for Norway, reported in Fagereng, Guiso, Malacrino, and Pistaferri (2020); and for Denmark, reported in Boserup, Kopczuk, and Kreiner (2018). The R^2 for the U.S. in Charles and Hurst (2003), who use a sample of parent-child pairs with positive wealth in both generations, are slightly higher.

Here, however we are attempting to explain 25-year long histories of individual wealth holdings. Thus, we view this comparable magnitude as evidence for the success of our procedure. Moreover, we take the explanatory power of ex ante variables as showing that there is substantial variation in outcomes later in life that cannot be captured by initial characteristics. This result is consistent with the existence of many ex post factors (such as savings decisions, return risk, and labor market risk) that drive wealth accumulation but are not predetermined at age 30.

Finally, it is worth noting that the relatively low explanatory power of observables does not appear limited to our application to wealth mobility. For instance, Ahn, Hobijn, and Şahin (2023), who recover latent worker groups from individual labor market histories, also find that observable demographic characteristics have limited explanatory power for group membership.

E.4. Explanatory power of additional covariates

Finally, we reproduce the exercise above, computing the distance-weighted and unweighted classification rates of the estimated multinomial logit model (equation 6) with the additional covariates described in Appendix E.2.

We present the results in Table E.2. The distanced-weighted classification rate increases up to 20 percent, with the introduction of the individuals' initial wealth ventile accounting for 15 percentage points of the total classification rate. A similar increase takes place for the unweighted classification rate that increases to 10.6 percent (over the random classification rate of 29.3 percent). The individuals' initial wealth ventile accounts for 7.9 percentage points of the total classification rate. As we discussed in Appendix E.2, we see the large role of initial wealth as being consistent with the patterns of segmented wealth mobility we document.

Importantly, while the partial contribution of parental wealth and education decreases, these changes are moderate, particularly with respect to the large increase in explanatory power afforded by the inclusion of the individuals' initial position in the wealth distribution. The (level) contribution of parental wealth and education to the distance-weighted classification rate go down 33 and 14 percent, respectively, while the overall explanatory power goes up 240 percent.

TABLE E.2. Predictive power of ex ante characteristics with additional covariates

Total Contribution	Partial Contribution					
	Par. Wealth	Edu.	Sex & Birthplace	Par. Edu.	Par. Bus.	Own State
Share of Distance Variation Explained by Variable (pp)						
20.0	8.1	9.9	2.8	1.5	3.0	74.6
Share of Individuals Correctly Classified (pp)						
10.6	7.2	10.7	4.0	1.7	2.5	73.9

Notes: The table reports the distance-weighted and unweighted probability of belonging to an individual's true group relative to a naive random classification with an unweighted classification rate of 29.3 percent. The distance corresponds to the measure in equation E.1. The combined explanatory power of all covariates is reported in column "Total Contribution." The remaining columns report the partial contribution of each variable category as a share (in percentage points) of their combined explanatory power in column "Total Contribution." The classification model corresponds to our estimated multinomial logit presented in equation (6) with the additional covariates introduced in Appendix E.2. The partial contribution of each variable category is obtained through the Shapley-Owen decomposition, which averages across permutations of decompositions and sums to the total contribution, as described in Appendix F.

Appendix F. The Shapley-Owen-Shorrocks Decomposition

Given an arbitrary function $Y = f(X_1, X_2, ..., X_n)$, the Shapley-Owen-Shorrocks decomposition is a method to decompose the value of $f(\cdot)$ into each of its arguments $X_1, X_2, ..., X_n$. Intuitively, the contribution of each argument if it were to be *"removed"* from the function. However, because the function can be nonlinear, the order in which the arguments are removed matters in general for the decomposition. The function f can be the outcome of a regression, like the predicted values or sum of square residuals, or the output of a structural model, such as a counterfactual value for a variable given a list of model parameters or components, or a transformation of the sample, for example the Gini coefficient.

The Shapley-Owen-Shorrocks decomposition is the unique decomposition satisfying two important properties. First, the decomposition is exact decomposition under addition, letting C_j denote the contribution of argument X_j to the value of the function $f(\cdot)$,

$$\sum_{j=1}^{n} C_j = f(X_1, X_2, ..., X_n), \tag{F.1}$$

so that $C_j/f(\cdot)$ can be interpreted as the proportion of $f(.)$ that can be attributed to X_j.[68]

[68]The interpretation holds as long as f is non-negative. If f can take negative values, then the interpretation of C_j under the exact additive rule can be misleading as some arguments can have $C_j < 0$.

Second, the decomposition is symmetric with respect to the order of the arguments. That is, the order in which the variable X_j is removed from $f(\cdot)$ does not alter the value of C_j.

The decomposition that satisfies both those properties is

$$C_j = \sum_{k=0}^{n-1} \frac{(n-k-1)!k!}{n!} \left(\sum_{s \subseteq S_k \setminus \{X_j\} : |s|=k} \Big[f(s \cup X_j) - f(s) \Big] \right), \tag{F.2}$$

where n is the total number of arguments in the original function f, $S_k \setminus \{X_j\}$ is the set of all *"submodels"* that contain k arguments and exclude argument X_j.[69] For example,

$$S_{n-1} \setminus X_n = f(X_1, X_2, ..., X_{n-1})$$
$$S_1 \setminus X_n = \{f(X_1), f(X_2), ..., f(X_{n-1})\}.$$

The decomposition in (F.2) accounts for all possible permutations of the decomposition order. Thus, $\frac{(n-k-1)!k!}{n!}$ can be interpreted as the probability that one of the particular submodel with k variables is randomly selected when all model sizes are all equally likely. For example, if $n = 3$, there are submodels of size $\{0, 1, 2\}$. In particular, there are 2^2 permutation of models that exclude each variable: $\{\underbrace{(0,0)}_{k=0}, \underbrace{(1,0), (0,1)}_{k=1}, \underbrace{(1,1)}_{k=2}\}$.

$$k = 0 : \frac{(n-k-1)!k!}{n!} = \frac{(3-0-1)!0!}{3!} = \frac{1}{3}$$
$$k = 1 : \frac{(n-k-1)!k!}{n!} = \frac{(3-1-1)!1!}{3!} = \frac{1}{6}$$
$$k = 2 : \frac{(n-k-1)!k!}{n!} = \frac{(3-2-1)!2!}{3!} = \frac{1}{3}$$

Nonlinear example

We illustrate the value of this decomposition with a simple nonlinear model including $n = 3$ variables:

$$Y = f(X_1, X_2, X_3) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 X_2. \tag{F.3}$$

[69] We abuse notation here. A submodel is an evaluation of function f with only some of its arguments. This language is motivated by the function corresponding in practice to the outcome of a regression or structural model. Formally, when we write $f(X_1)$, we mean $f(X_1, \emptyset_2, ..., \emptyset_n)$, where we assume the j-th argument of the function can always take on a null value denoted $\emptyset_j$. In our regression example below, this null value corresponds to a zero valued regressor or parameter. In the case of the structural model, this null value can correspond to setting some parameters to a predetermined value or excluding certain model components, like the adjustment of prices or a specific shock agents face.

The objective is to decompose the value of Y into the contribution (or partial effect) of each variable.

Removing X_1

There are four possible models that exclude X_1—one with no variable, two with one variable, and one with two variables:

$$\begin{aligned} k = 0 &: \beta_0 \\ k = 1 &: \{\beta_0 + \beta_2 X_2\ ,\ \beta_0\} \\ k = 2 &: \beta_0 + \beta_2 X_2 + \beta_3 X_3 X_2 \end{aligned}$$

In all four models, the partial effect of including X_1 is always $f(s \cup X_1) - f(s) = \beta_1 X_1$. This reflects the fact that the order in which variables are included does not matter to construct C_1:

$$C_1 = \sum_{k=0}^{2} \frac{(3-k-1)!k!}{3!} \left(\sum_{s \subseteq S_k \setminus \{X_3\} : |s| = k} \Big[f(s \cup X_j) - f(s) \Big] \right) = \beta_1 X_1 \tag{F.4}$$

This would be the same for any argument X_j entering linearly into f an arbitrary number of variables: $Y = f(X_1, X_2, X_3, X_4, \ldots, X_n) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 X_2 + \sum_{j=4}^{n} \beta_j X_j$. The only difference is that the number of submodels grows exponentially, 2^{n-1}, but the partial effect of including X_j for some $j \in \{4, \ldots, n\}$ is always $C_j = \beta_j X_j$.

Removing X_2

In this case, the partial effect can be decomposed into all the possible ways X_2 can be added into the model, $f(s \cup X_2) - f(s)$, these are

$$\begin{aligned} k = 0\ (\emptyset_1, \emptyset_3) &: \beta_0 + \beta_2 X_2 - \beta_0 = \beta_2 X_2 \\ k = 1\ (X_1, \emptyset_3) &: \beta_0 + \beta_1 X_1 + \beta_2 X_2 - (\beta_0 + \beta_1 X_1) = \beta_2 X_2 \\ k = 1\ (\emptyset_1, X_3) &: \beta_0 + \beta_2 X_2 + \beta_3 X_2 X_3 - \beta_0 = \beta_2 X_2 + \beta_3 X_2 X_3 \\ k = 2\ (X_1, X3) &: \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_2 X_3 - (\beta_0 + \beta_1 X_1) = \beta_2 X_2 + \beta_3 X_2 X_3 \end{aligned}$$

Here, the partial effects of adding X_2 are not the same across submodels because X_2 enters nonlinearly into the original model. The symmetric property of the decomposition takes care of this.

$$\begin{aligned} C_2 &= \underbrace{\frac{1}{3}\beta_2 X_2}_{k=0} + \underbrace{\frac{1}{6}(\beta_2 X_2) + \frac{1}{6}(\beta_2 X_2 + \beta_3 X_2 X_3)}_{k=1} + \underbrace{\frac{1}{3}(\beta_2 X_2 + \beta_3 X_2 X_3)}_{k=2} \\ &= \beta_2 X_2 + \frac{1}{2}\beta_3 X_2 X_3 \end{aligned} \tag{F.5}$$

The result is quite intuitive. $\beta_2 X_2$ appears in all submodels; hence, its probability of appearing in the decomposition is 1. $\beta_3 X_2 X_3$ appears in two of the four submodels;

hence, its probability of appearing is 1/2. Weighting each term by its probability of appearing in the decomposition ensures symmetry.

Removing X_3

We proceed in the same way for X_3 as we did for X_2. There are four submodels. In two of them, the effect of adding X_3 is null, because X_2 is not in the model. In the two remaining submodels, the effect is $\beta_3 X_2 X_3$. Hence,

$$C_3 = \frac{1}{2}\beta_3 X_2 X_3. \tag{F.6}$$

Finally, we verify the decomposition:

$$\begin{aligned} C_1 + C_2 + C_3 &= \beta_1 X_1 + \left(\beta_2 X_2 + \frac{1}{2}\beta_3 X_2 X_3\right) + \left(\frac{1}{2}\beta_3 X_2 X_3\right) \\ &= \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_2 X_3 \\ &= f(X_1, X_2, X_3) - \beta_0 \\ &= f(X_1, X_2, X_3) - f(\emptyset_1, \emptyset_2, \emptyset_3). \end{aligned}$$

Note: The decomposition is additive with respect to the reference "null" model where none of the variables is included. This is made apparent in the previous result, where the decomposition does not include the value of β_0.

R-Squared

Finally, we consider a decomposition of the coefficient of determination in the linear model. Our use of the decomposition applies this for a nonlinear model (combining the insights from this and the preceding example).

Consider a linear regression model with n regressors and $i = 1, \ldots, M$ observations,

$$y_i = \mathbf{x}_i'\beta + u_i = \beta_0 + \sum_{j=1}^{n} \beta_j x_{ij} + u_i, \tag{F.7}$$

and define the average value of y as $\overline{y} \equiv \sum_{i=1}^{M} y_i/M$ and the predicted value

$$\hat{y}_i = \mathbf{x}_i'\hat{\beta} = \hat{\beta}_0 + \sum_{j=1}^{n} \hat{\beta}_j x_{ij}, \tag{F.8}$$

where we assume that all regressors have zero mean so that $\hat{\beta}_0 = \overline{y}$.

The function of interest is $f(X_1, \ldots, X_K) = R^2$, defined as the explained sum of

squares *SSE* over the total sum of squares *SST*

$$R^2(X_1, X_2, .., X_n) = \frac{SSE}{SST} = \frac{\sum_{i=1}^{M}(\hat{y}_i - \overline{y})^2}{\sum_{i=1}^{M}(y_i - \overline{y})^2}. \tag{F.9}$$

This makes it clear that the function being decomposed is nonlinear even though the model that generates it is itself linear.

Note: The reference value for the R^2 in the Shapley-Owen-Shorrocks decomposition is given by the model without regressors, satisfying

$$R^2(\emptyset) = \frac{\sum_{i}^{M}(\hat{\beta}_0 - \overline{y})^2}{\sum_{i}^{M}(y_i - \overline{y})^2} = 0, \tag{F.10}$$

so that, in this case, the decomposition recovers the level of the R^2 of the full model (with all variables), unlike the previous example.

Details of the decomposition when n = 3 Consistent with the previous example, we show the decomposition for $n = 3$ regressors. As before, we abuse notation by only listing the arguments being included in each submodel. The contribution of each variable is:

$$\begin{aligned} R_1^2 = {} & \frac{1}{3}\left[R^2(X_1) - R^2(\emptyset)\right] + \frac{1}{6}\left(\left[R^2(X_1, X_2) - R^2(X_2)\right] + \left[R^2(X_1, X_3) - R^2(X_3)\right]\right) \\ & + \frac{1}{3}\left[R^2(X_1, X_2, X_3) - R^2(X_2, X_3)\right]; \end{aligned} \tag{F.11}$$

$$\begin{aligned} R_2^2 = {} & \frac{1}{3}\left[R^2(X_2) - R^2(\emptyset)\right] + \frac{1}{6}\left(\left[R^2(X_1, X_2) - R^2(X_1)\right] + \left[R^2(X_2, X_3) - R^2(X_3)\right]\right) \\ & + \frac{1}{3}\left[R^2(X_1, X_2, X_3) - R^2(X_1, X_3)\right]; \end{aligned} \tag{F.12}$$

$$\begin{aligned} R_3^2 = {} & \frac{1}{3}\left[R^2(X_3) - R^2(\emptyset)\right] + \frac{1}{6}\left(\left[R^2(X_3, X_2) - R^2(X_2)\right] + \left[R^2(X_1, X_3) - R^2(X_1)\right]\right) \\ & + \frac{1}{3}\left[R^2(X_1, X_2, X_3) - R^2(X_2, X_1)\right]. \end{aligned} \tag{F.13}$$

Summing across all the contributions we obtain back $R^2(X_1, X_2, X_3)$,

$$R_1^2 + R_2^2 + R_3^2 = R^2 = f(X_1, X_2, X_3). \tag{F.14}$$

Note: The value of the contribution differs from the standard definition of partial R-squared. This is because the partial R-squared is an all-else-being-equal comparison of excluding regressor X_j from the regression. It does not satisfy the exact decomposition requirement or (when applied iteratively) the symmetry requirement.