



Canadian Labour Economics Forum

WORKING PAPER SERIES

The Life-Cycle Dynamics of Wealth Mobility

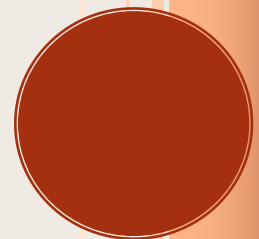
Richard Audoly (Federal Reserve Bank of
New York)

Rory McGee (Western University)

Sergio Ocampo (Western University)

Gonzalo Paz-Pardo (European Central Bank)

CLEF WP #68



The Life-Cycle Dynamics of Wealth Mobility*

Richard Audoly Rory M^cGee Sergio Ocampo Gonzalo Paz-Pardo

April 2024

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. Movements are tied to differential income trajectories and business activities across groups. 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.

M^cGee: 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, Juan Herrero, Joachim Hubmer, John Bailey Jones, Bariş Kaymak, Hannes Malmberg, Elena Manresa, Cormac O'Dea, Serdar Ozkan, 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. M^cGee 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 that 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.

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-fourth of individuals are in the same quintile of the distribution after 25 years. 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 full realized trajectories. We group 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

¹Low 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).

²See 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 to at least to Wedgwood (1929).

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 to 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, with one rising through the distribution and the other falling. 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 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 (that has a larger share of capital income). By contrast, fallers have similar, but ultimately less successful, entrepreneurial activities compared to those at the top. The individuals at the bottom of the distribution are very different from the others: their incomes are persistently lower, they stay renters throughout their lives, and they rarely own businesses.

Turning to the composition of our main groups, we look into their most

³Differences in realized wealth rank trajectories correspond to meaningful differences in levels reflecting the high degree of wealth inequality in Norway. For instance, the gap at age 55 between the two groups in the middle of the distribution represents a difference in their net worth of almost 600,000 US dollars.

representative subgroups exploiting the hierarchical nature of our clustering algorithm. Those rising mostly differ in the timing of their movements, with subgroups experiencing similar gains in relative wealth positions but at different ages, a pattern we relate to their educational attainment. On average, the risers are in the 40th percentile of the wealth distribution around age 30 and climb to the 70th percentile by age 55. There is substantially more heterogeneity in the internal patterns of those falling through the distribution. Two subgroups are continuously falling; the one experiencing the largest fall goes from the 75th to the 38th percentile between ages 30 and 55, a pattern we tie to declining business performance. A third group experiences a rapid rise by age 45 but later falls back down; these movements coincide with high early marriage rates and late divorce rates. We explore the heterogeneity within all major groups in Section 6.

Finally, we contrast the role of individuals' circumstances, including parental wealth and education, in predicting full wealth rank histories.⁴ 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 part of the group that is persistently at or near the top of their own generation's distribution, compared to those born to parents at the bottom of the distribution. In contrast, those born to parents at the bottom of the distribution are not only likely to be poorer at a given age; they are also more likely to be persistently poor throughout their lives.

However, parental wealth plays a more limited role for individuals who experience a rise or fall through the distribution. For these individuals, 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

⁴Our 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.

their parental background, those without post-secondary education are between 5 and 10 percentage points more likely to be fallers than those with at least undergraduate degrees. Overall, parental wealth and human capital each account for 40 percent of the explained variation in group membership—highlighting the importance of hereditary advantage in wealth dynamics (Becker and Tomes 1979).⁵

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 parent’s ranks as they age. Not only does this mirror the intragenerational mobility trend we document, but we find that the same individuals drive both population trends. As risers 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.⁶ 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.

⁵Our 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.

⁶See 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 behaviour differs over a decade by race, education, household demographics, and initial wealth using the Panel Study of Income Dynamics. We, instead, study mobility with wealth trajectories over 25 years without conditioning on specific variables.

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](#); and [Fagereng, Mogstad, and Rønning 2021](#)). Our methodology allows us to go beyond comparing across generations at a given point in their life cycle by considering the full wealth histories of individuals. We show that both parental background and the individuals' positions in the wealth distribution near the beginning of their work-life have long-lasting impacts on the wealth trajectories of individuals. 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, 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 8 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.⁷ We report monetary values in 2019 US dollars.

The coverage and properties of the Norwegian administrative datasets is 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. We start by highlighting the key strengths of our data.

First, Norway has recorded wealth in its tax returns since 1993, providing us with a long panel with twenty-five years of observations. This long panel allows us to track individuals over important phases of their life cycles. Tracking individuals is crucial to understand mobility over long horizons and to differentiate the life-cycle trajectories experienced by individuals, as we do when we document the trajectories of wealth mobility using our clustering procedure.

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

⁷The 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 \(2015\)](#) provide a detailed discussion of income tax records.

capture in survey data.⁸ 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 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. 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. Using individual level data allows us to consider how mobility varies with sex as well as to track individuals through household formation and dissolution. In Section 6.4, we show that these dynamics are related to the mobility patterns of subgroups of individuals.

We also observe the value of various asset classes included in individuals' wealth tax returns. However, the returns do not include transactions within classes. The single largest asset for most individuals is housing. We adjust housing values using the reported values in [Fagereng, Holm, and Torstensen \(2020\)](#) treating condominiums and other properties separately.⁹ We aggregate primary residences, secondary residences and leisure properties, and foreign residences into a property-asset class. Finally, we define

⁸This 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.

⁹We thank [Fagereng, Holm, and Torstensen](#) for providing us with updated adjustment values covering our sample period.

a home ownership indicator excluding secondary and foreign properties.

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.¹⁰ A person'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.

Two types of assets are missing from our data. First, assets individuals choose to 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. Second, we lack information on pension entitlements, including employer-provided pension plans. Private pensions represent less than 20 percent of all pensions in Norway, while pay-as-you-go public pensions makes up the remaining 80 percent (Ozkan, Hubmer, Salgado, and Halvorsen 2023).¹¹

2.2. Additional data: income and demographics

We also use high-quality information on individual incomes from income tax records, analogous to the wealth tax records described in the preceding section. 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.

¹⁰We 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.

¹¹Pay-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. See Fagereng, Holm, Moll, and Natvik (2019, Appendix C.6) for details of the Norwegian Public Pension system.

We also have access to detailed information on individual education levels and fields of study, according to the Norwegian Standard Classification of Education (NUS2000).¹² 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). This lets us classify individuals as cohabiting (or married), divorced, and widowed.

2.3. Sample selection

We begin with the universe of Norwegian tax residents at any point 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 represents our main object of interest.¹³ This birth cohort is first observed in its early thirties in our data (1993–2017) and is therefore observed for a significant fraction of their work lives. In addition, this cohort is not affected by the changes in the compulsory school age taking place in 1959. This reform was not implemented uniformly across place and time; see [Black, Devereux, and](#)

¹²See <https://www.ssb.no/en/klasse/klassifikasjoner/36>.

¹³To illustrate the value of our Norwegian administrative 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 level of wealth is only observed consistently since 1999, implying an average of just six consecutive wealth 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.

Salvanes 2005; and Bhuller, Mogstad, and Salvanes 2017 for more details.

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. Our balancing eliminates attrition driven by migration and mortality. Given the age range we consider, increasing mortality in late middle age drives a large share of this sample selection criteria.

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 begin by measuring the persistence of individuals' rank in the wealth distribution as they age.

3.1. Wealth ranks and the Norwegian distribution of wealth

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

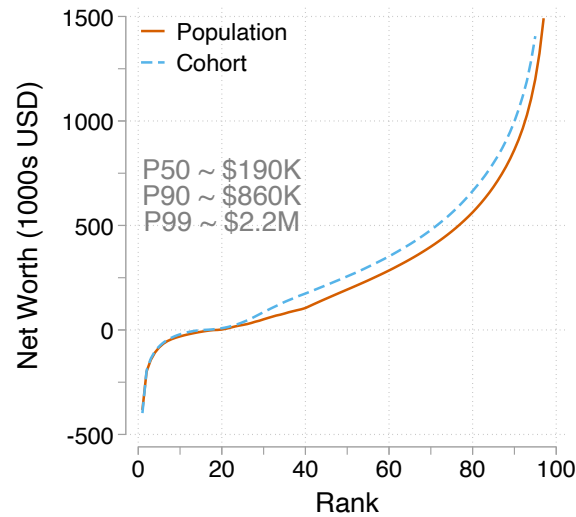
$$y_{i,t} = 100 \times F_w \left(w_{i,t} \mid t, i \in \text{Birth Cohort: 1960–64} \right), \quad (1)$$

where F_w denotes the empirical cumulative distribution of wealth. We multiply by 100 to express ranks on a percentile scale. Crucially, all comparisons use other individuals of the same cohort as the reference group. As a result, our rank measure is not affected by cross-cohort or cross-age comparisons.¹⁴

Ranks have several attractive properties over alternative monotonic transformations

¹⁴Importantly, doing this also purges ranks from time effects varying by age. For instance, all members of our sample experience the effects of the 2008 global recession at approximately the same age. Therefore, we do not consider cross-cohort differences in patterns of life-cycle wealth accumulation (see, for example, Gale, Gelfond, Fichtner, and Harris 2021; and Paz-Pardo 2024, who document these changes in the U.S.).

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 average wealth holdings in 2019 US dollars by percentile.

of net worth. First, they are well-defined for individuals with negative or zero net worth. Second, they are easy to interpret in terms of relative mobility in either levels or changes. Third, because they compress the right tail of the distribution, they capture diminishing marginal gains from wealth for individuals. 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](#)).¹⁵

In Figure 1, we report wealth in US dollars by rank in the wealth distribution. As in most advanced economies, wealth in Norway is very unequally distributed. For reference, the 90th percentile of wealth in Norway is close to 860,000 US dollars, while it is 620,000 dollars in the U.S. ([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 is equivalent to going from 190,000 to 250,000 US dollars of

¹⁵We also report results for trajectories of (log) wealth and for alternative measures of the position of individuals in the wealth distribution, such as their position in the Lorenz curve, which deliver similar qualitative findings to our main analysis. These results are described in 8.2.

net worth. The only part of the distribution in which rank changes do not translate into substantial movements in wealth levels is the narrow window around zero wealth (15th–20th percentile). Moreover, changes in rank reflect meaningful differences in wealth even at younger ages, when our cohort is less wealthy than the population at large (see Appendix A).

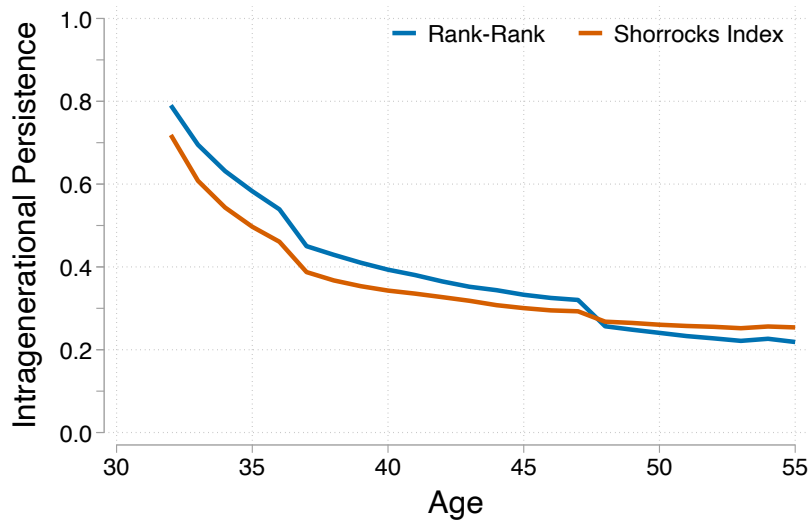
3.2. Rising wealth mobility along the life cycle

We now measure intragenerational wealth mobility: to what extent do individuals transition across the wealth distribution over their lives? We measure life-cycle wealth mobility with the persistence of individuals' rank in their cohort's wealth distribution. This persistence summarizes the evolution of the histories of wealth ranks across the population. For instance, these histories indicate high persistence (and low mobility) when individuals tend to remain in the same relative positions within their cohort throughout their lives.

We use two measures of rank persistence. First, we use the auto-correlation of ranks as individuals age with respect to their initial position in the 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, $y_{i,1993}$.¹⁶ Second, we use the [Shorrocks \(1978\)](#) index of mobility given by the share of individuals who remain in their initial quintile of the wealth distribution as they age. This provides age-varying measures for the persistence of individuals' position in the wealth distribution. For both measures, a persistence of one indicates that there is no mobility, and individuals do not change their position relative to the wealth of others. A lower persistence implies more mobility in the individuals' rank (or quintile).

¹⁶Formally, we estimate $y_{i,t} = \alpha_t + \rho_t y_{i,1993} + u_{i,t}$. [Chetverikov and Wilhelm \(2023\)](#) derive the asymptotic distribution of the rank-rank OLS slope under several alternative assumptions, such as the presence or absence of ties in the ranked variables or additional covariates.

FIGURE 2. Intragenerational persistence of wealth ranks



Notes: Figure 2 plots the rank-rank and the Shorrocks' (1978) persistence measures for intra-generational mobility. 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 Shorrocks index corresponds to the share of individuals who at time t are in the same wealth quintile that they were in 1993.

Figure 2 shows that both intragenerational persistence measures decrease as individuals grow older, evidencing rising wealth mobility along the life cycle. Thus, individuals experience increasingly large (cumulative) changes in rank that persistently change their position in the wealth distribution. Most of these changes take place at relatively younger ages, when persistence declines rapidly, and are less frequent from age 37 onward, when persistence stabilizes around 0.3 before decreasing again a decade later.¹⁷

Although these results establish a trend of increasing cumulative intragenerational mobility, they remain silent over how broad-based this mobility is, as well as what (or who) is driving the trend. The persistence measures shown in Figure 2 collapse the

¹⁷In Appendix D, we compute the same measures of persistence for intergenerational mobility, focusing on the relationship between individuals' current wealth rank and their parents' rank in 1993. We find that, in contrast to the results in Figure 2, there is a trend of decreasing wealth mobility along the life cycle; the relationship between an individuals' positions in the wealth distribution and those of their parents grows stronger as individuals age.

myriad of wealth trajectories experienced by individuals through their lives into a single aggregate time series; they necessarily obscure the heterogeneity in trajectories across the population. Some of these trajectories correspond to individuals with relatively stable ranks (and low mobility) and some to individuals who undergo large changes, rising or falling through the wealth distribution. Put another way, the trends in Figure 2 tell us that rank changes occur, but they contain no information about the shapes of individuals' typical wealth histories, how frequent stable or changing trajectories are, or what the usual timing and magnitude of rank changes are.

To understand the underlying life-cycle patterns of wealth behind increasing mobility, and the economic mechanisms that shape them, we move to analyze the distribution of wealth trajectories. We do this by decomposing 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 also capture persistent differences in mobility, describing a pattern of segmented wealth mobility.

4. Grouping life-cycle trajectories of mobility

Our key object of interest is the degree and timing of wealth mobility over the life cycle of individuals, which we measure using the distribution of wealth rank histories. In this section, we describe our empirical approach to studying this high-dimensional object: we use hierarchical agglomerative clustering to group individuals and recover a limited number of typical wealth trajectories.¹⁸ 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

¹⁸It is possible to include other outcomes or covariates that differ across groups, such as their income or portfolio composition. However, in practice, doing this may introduce more noise than additional information; accordingly, we focus on wealth ranks in our main analysis.

the panel is described by the vector of ranks:

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

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 we therefore proceed by reducing this object to a small number of groups. We recover a set of $G > 1$ disjoint groups (or clusters) of individuals, so that each individual i is assigned to one of these groups, $g_i \in \{1, \dots, G\}$. 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. 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:¹⁹

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

where g and g' are disjoint groups, N_g denotes the number of observations in group g ,

¹⁹Alternative 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>.

and $\bar{\mathbf{Y}}^g$ is the centroid (average) of the observations in group g .²⁰

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.

Selecting the number of groups. A key feature of our hierarchical clustering algorithm, 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.²¹ Instead, implementing the algorithm recovers a complete hierarchy of nested groups. This makes it easy 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.

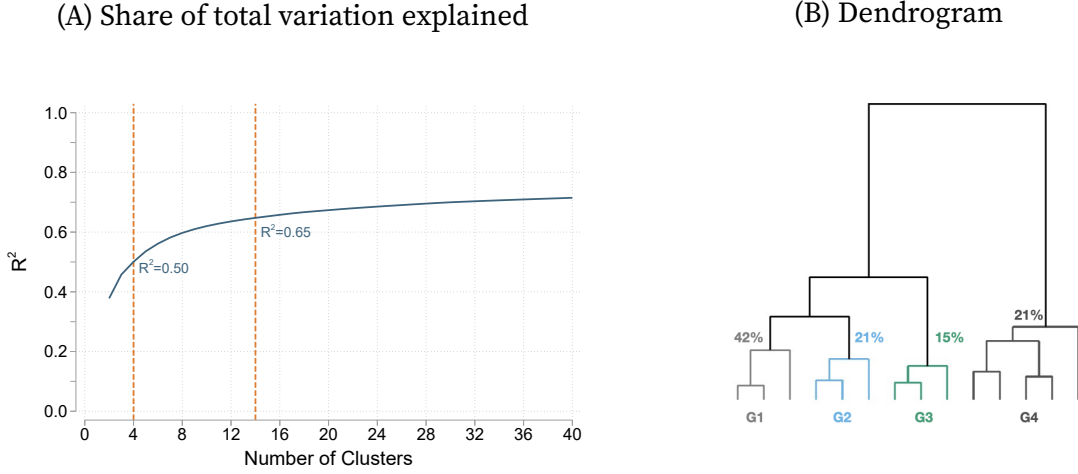
Selecting the baseline number of groups used in our analysis requires that we trade off two objectives: (i) having enough groups to represent the distribution of wealth rank histories, and (ii) having a parsimonious description of trajectories. Fewer groups provide a more easily interpretable picture of wealth mobility but may obscure relevant variation in the trajectories of the group’s members.

We 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 $\mathcal{G}_G = \{g_i\}_{i=1}^G$ over the set of

²⁰Given a set of underlying groups G^* , 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 $G < G^*$ groups. [Borysov, Hannig, and Marron \(2014\)](#) show this is the case for Ward’s method as either $T/N \rightarrow \infty$, with T growing faster than N , or only $T \rightarrow \infty$, 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.

²¹We discuss these differences and report results under alternative clustering approaches in Section 8.

FIGURE 3. Choice of number of groups



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.

individuals with $G > 1$ groups, this measure is

$$R^2 = 1 - \frac{\sum_{i,t} \left(y_{i,t} - \bar{y}_t^{g(i)} \right)^2}{\sum_{i,t} \left(y_{i,t} - \bar{y} \right)^2}, \quad (4)$$

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 3A presents the R^2 for the partitions produced by our hierarchical clustering algorithm for $G = 1, \dots, 40$. 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 3B) only

increases the R^2 from 50 percent to 65 percent.²² Using four groups therefore captures typical wealth accumulation trajectories well, given our objective to summarize the joint distribution of wealth over an individual’s work life, and we select them as our baseline.

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 3B. 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 6.4 and in Appendix B.4.

5. Segmented wealth mobility

We group individuals who experience similar trajectories through the wealth distribution as they age, using agglomerative hierarchical clustering. The resulting typical trajectories of wealth experienced by individuals in our four baseline groups imply very different mobility patterns. We show that these patterns are the results of segmented mobility among the individuals in our sample, with two of the groups driving the trend of increasing intragenerational mobility described in Figure 2.

²²The R^2 defined in (4) is a function of the variation between groups, captured by the *between* R^2 , and within groups, captured by the *within* R^2 . 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. 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. The within R^2 for $G = 4$ groups is close to 15 percent, but it more than doubles for $G = 14$ groups. We discuss these measures in Appendix B.1.

5.1. Typical wealth trajectories over the life cycle

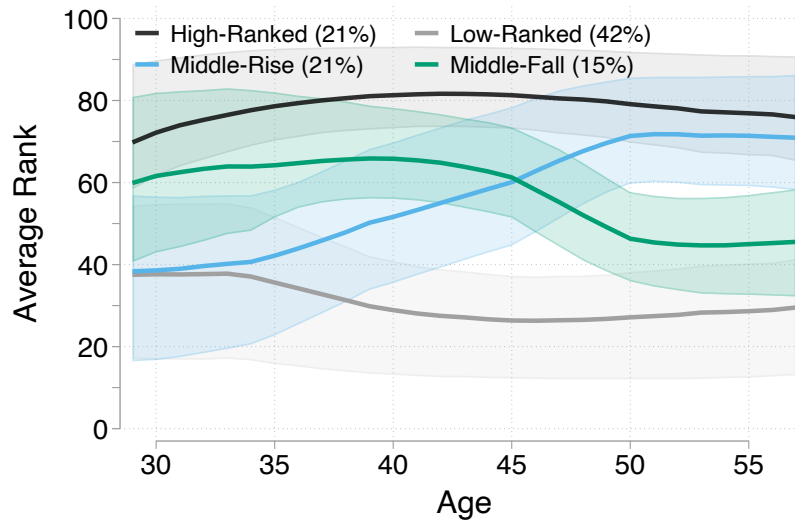
We begin by reporting the typical wealth rank trajectories of the individuals in each of our four main groups in Figure 4. The typical trajectories have groups remaining at the bottom, in the middle, and at the top of the distribution across their life cycle, with the groups in the middle exhibiting 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*.²³

Two groups of individuals, which we label “high-ranked” and “low-ranked,” start their lives at the top or the 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 6.4) but that it tends to stay within the upper or lower part of the wealth distribution, as made clear by the small changes in the interquartile range of the distribution of ranks.

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 intragenerational mobility trend 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 limited to segments of the population.

²³The 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 8.

FIGURE 4. Life-cycle dynamics of wealth mobility

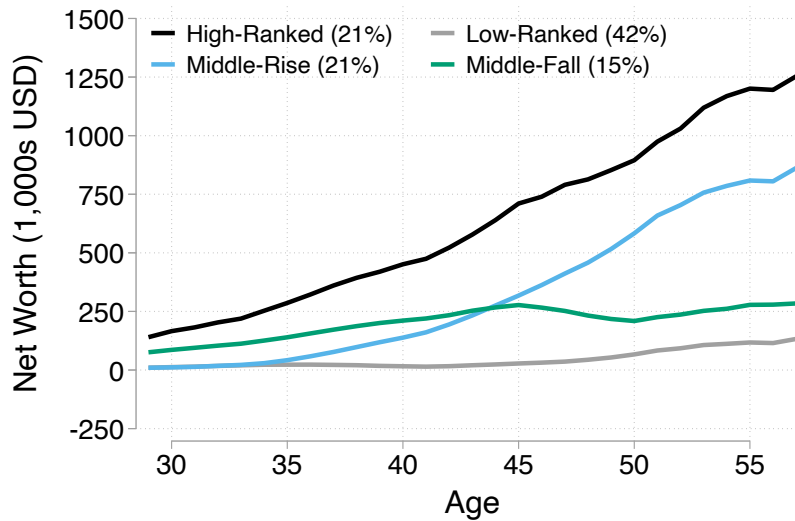


Notes: The figure plots the average wealth rank in each clustered group against the cohort's average age. The shaded areas correspond to the interquartile range of the rank distribution among the individuals of each group for each year. All individuals belong to the 1960–64 birth cohort. The clusters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

These typical wealth rank trajectories capture economically meaningful differences in the wealth trajectories of individuals, as we show in Figure 5. 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 5 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 very quickly as they age. Interestingly, the average rank among risers improves by accumulating net worth at a similar pace, but they start with zero wealth, on average, at age 30. By contrast, fallers have larger net worth than risers at

FIGURE 5. Average wealth by group



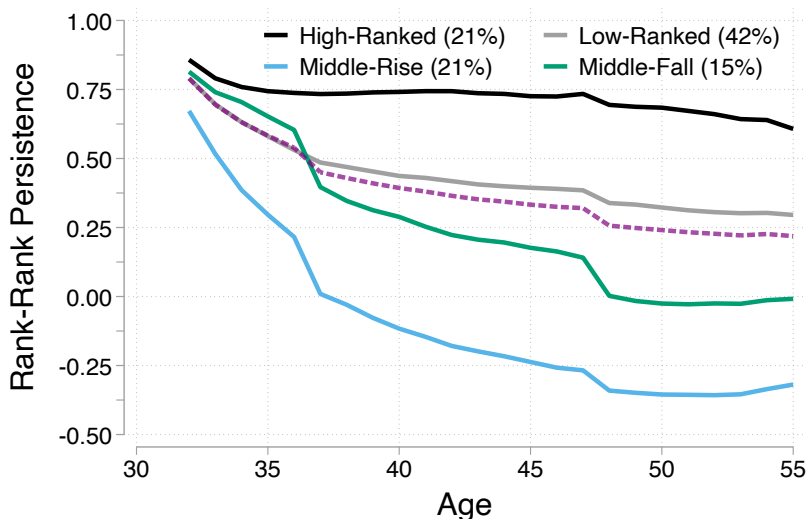
Notes: The figure plots the average wealth level in thousands of US dollars in each clustered group against the cohort's average age. All individuals belong to the 1960–64 birth cohort. The clusters are constructed from the balanced sample using hierarchical agglomerative clustering and Ward's method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

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 occurs for the 1960–64 birth cohort. 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 4 captures permanent differences in mobility across groups that, in turn, explain what (and who) is driving the trend of rising intragenerational mobility described in Section 3. To see this, we decompose the persistence of wealth ranks by computing the rank-rank persistence measure separately

FIGURE 6. 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, $y_{i,1993}$. To compute the rank-rank persistence for each group, we compute the auto-correlation from deviations from the cohort-wide average rank.

for each of our main four groups.²⁴ We present the results in Figure 6 and discuss the Shorrocks persistence measure in Appendix B.

Each group exhibits distinct patterns of wealth mobility. The high- 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 described earlier. The persistence of the relative wealth position of the 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 segment of the distribution, in part reflecting the larger size of this group.

²⁴Formally, we compute the heterogeneous rank-rank auto-correlation for each group g , ρ_t^g . We estimate $y_{i,t} = \alpha_t + \rho_t^{g(i)} y_{i,1993} + u_{i,t}$, where the intercept α_t does not depend on the group.

By contrast, mobility is higher for the groups of risers and fallers, reflecting large changes in ranks over the life cycle that result in low and even negative rank auto-correlation. 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. For this group, initial wealth rank is actually a negative predictor of future wealth ranks. For the group of fallers, persistence declines to zero, implying that their initial position in the wealth distribution has a very low correlation with their wealth rank at later stages in life.

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 and within groups

We now turn to exploring the ex post life-cycle characteristics of each of the four main wealth mobility groups, as well as the underlying heterogeneity revealed by their respective subgroups. 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 sources of income, entrepreneurship, portfolio composition, marriage and divorce, and inheritances. We choose these factors because they have been shown to be key determinants of wealth accumulation and wealth inequality in the literature (De Nardi and Fella 2017; Kuhn, Schularick, and Steins 2020; Hubmer, Krusell, and Smith 2021).

We find important differences across 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. Instead, we find that by dividing (and subdividing) the cohort, we identify more granular heterogeneity across and within groups. We summarize the resulting patterns next.

6.1. Portfolio composition

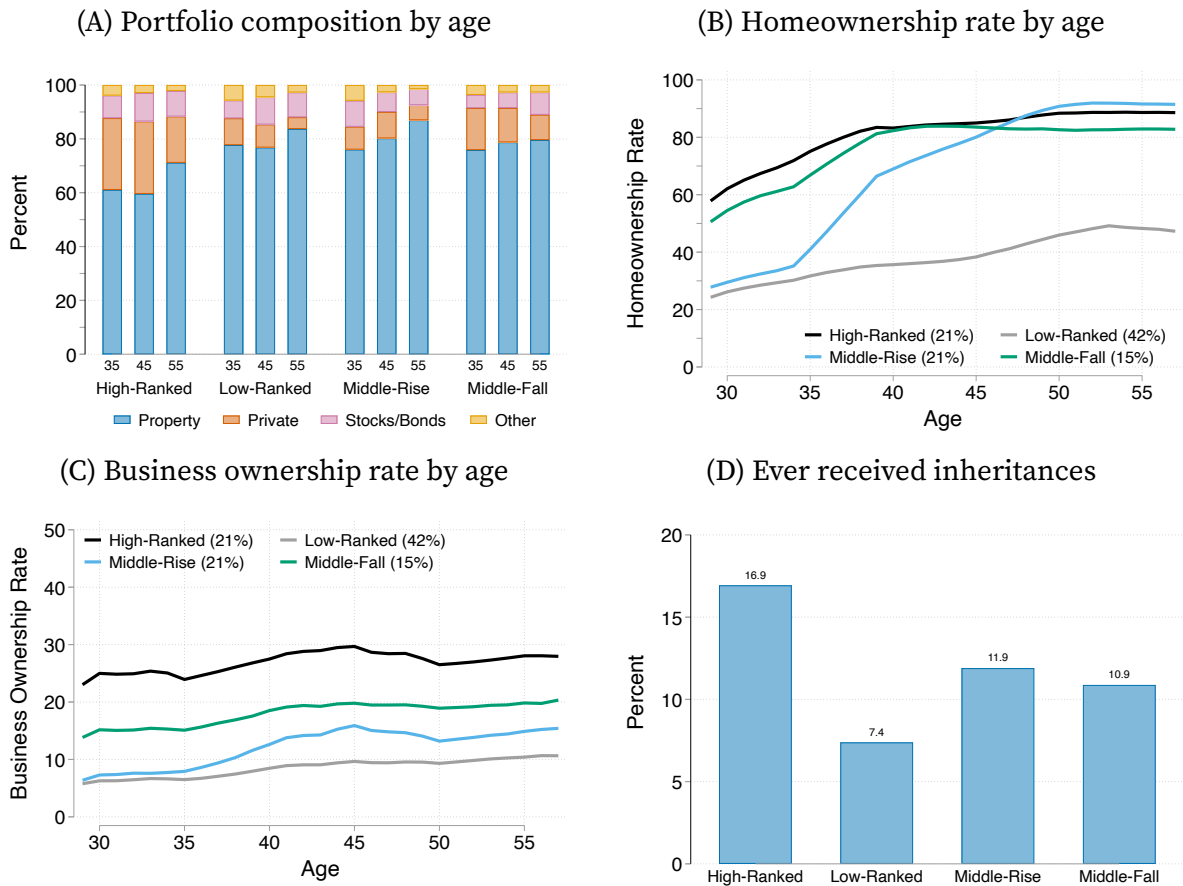
Figure 7 describes the portfolio composition of our four main groups. Panel A reports the share of assets accounted for by property, private business assets, financial assets (stocks, bonds, or bank accounts), and a residual category including vehicles and foreign assets.²⁵

Property represents the majority of household portfolios across all groups; its share increases slightly as homeownership rates increase. The increase in homeownership occurs in all groups as individuals age, but the age profiles differ markedly (Figure 7B). 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. As for the low-ranked groups, 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 shares of property assets in their portfolio and correspondingly higher shares of private business wealth, relative to the other groups. This pattern is driven by differences in

²⁵We find that the differences in net worth shown in Figure 5 mostly reflect differences in the accumulation of assets and not debt. Trajectories of debt show much lower dispersion and levels than the trajectories of assets, as we show in Figure B.3 in Appendix B.3.

FIGURE 7. Portfolio statistics by group



Notes: The figures present characteristics of the four main groups presented in Figure 4. 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.

the extensive margin of business operation (Figure 7C), as well as the group’s higher shares of capital income and self-employment income, which we discuss more later. 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.

Taken together, these patterns suggest that ownership and operation of profitable business assets are a crucial characteristics for many individuals who start and remain

at 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 out of labor income is instead the main driver for the latter group, in line with the findings of [Kuhn, Schularick, and Steins \(2020\)](#) for the U.S.

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. Figure 7D 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.²⁶

6.2. Income trajectories

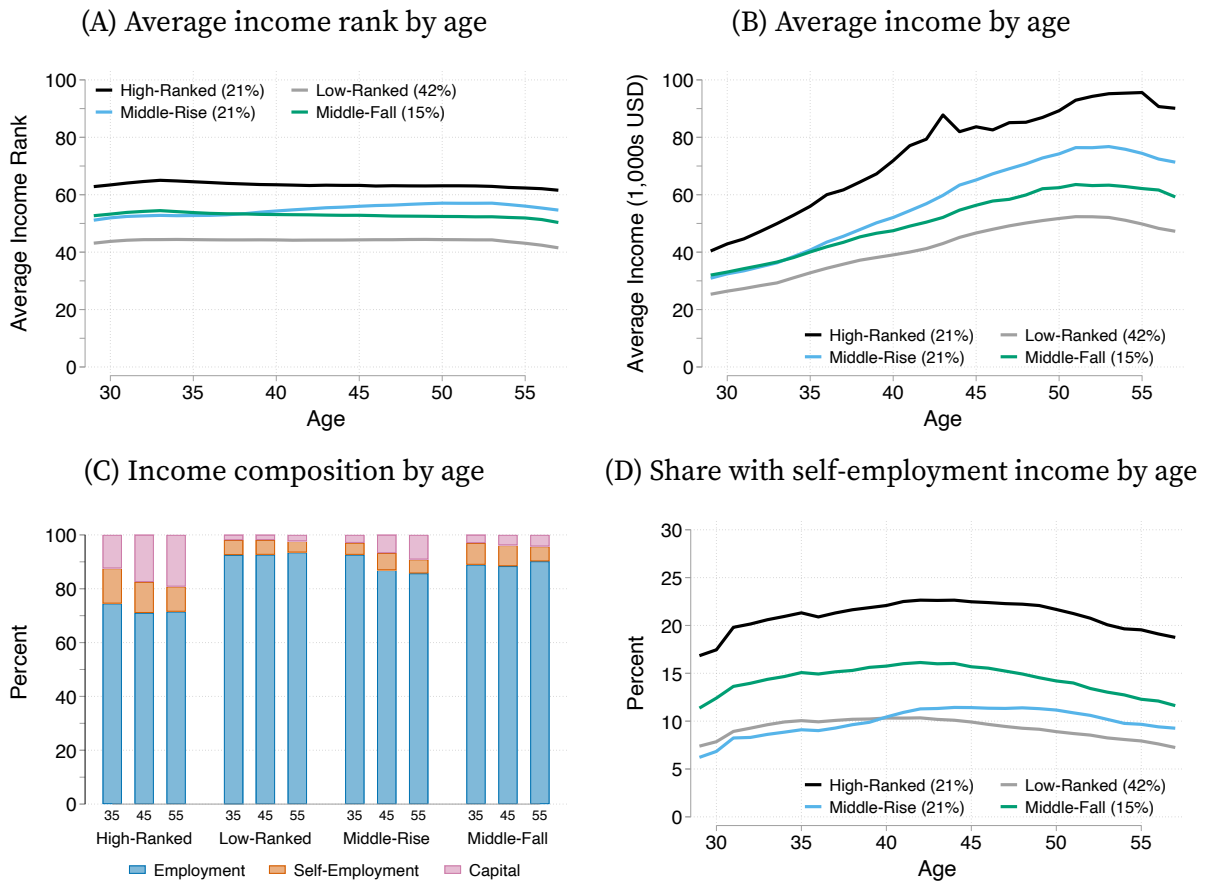
The income profiles of our main four groups are broadly consistent with the patterns of wealth mobility described previously, as we show in Figure 8. Nevertheless, the differences in income (levels and ranks) are smaller than the differences in wealth documented in Figures 4 and 5. Looking at income 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.²⁷

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.

²⁶Our 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. Consistent with the results we document below, they also find an important role for human capital.

²⁷The income rank gap between the high- and low-ranked groups translates into the income of those in the high-ranked group being, 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 8. 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 4. 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.

Fallers begin with more wealth than risers (Figure 5) 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 7 that these life-cycle differences in income are partly explained by higher educational attainment among risers (relative to all other groups), pointing to an important role for human capital accumulation for wealth mobility.

Figure 8C shows that the vast majority of income comes from wages and salaries

in all groups.²⁸ The high-ranked group is the only one for which capital income and self-employment income represent a sizeable part of their resources (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 8D), more than double than in 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

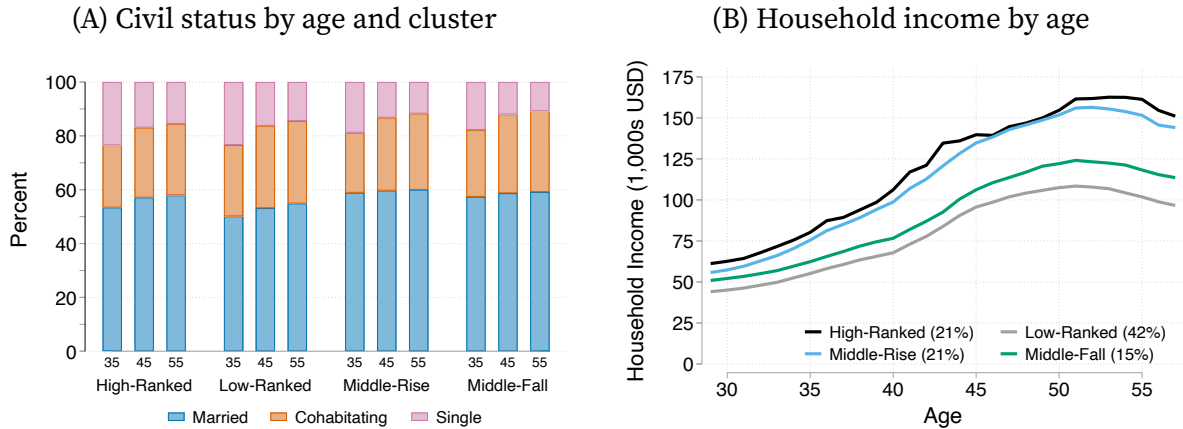
We also track the civil status of individuals in each group. Figure 9A reports the share of individuals who are married, cohabiting, and single at several points during the panel. Although individuals in the low-ranked group are slightly more likely to be single, we find relatively limited differences across groups. By age 45, over 80 percent of individuals in all groups are married or cohabiting.

The high rates of marriage and cohabitation make the trajectories of household income that we report in Figure 9B 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 equalize using Organization for Economic Co-operation and Development’s (OECD) equivalence scale based on the number of adults and children in the household. We also compute household wealth ranks, reported in Appendix B.3, and show that our main results are robust to using household wealth instead of individual wealth in Section 8.

We find that the life-cycle profiles of household income reinforce, rather than reduce, the differences in individual labor income across groups (Figure 8B). We find a large gap

²⁸We focus on sources of earned income, excluding income from public sector benefit programs. We show in Figure B.4 of Appendix B.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 receives unemployment benefits at age 40, and more than 20 percent receives disability benefits, compared to less than 5 and 10 percent, respectively, in the high-ranked group.

FIGURE 9. Household characteristics by group



Notes: The figures present characteristics of the four main groups in Figure 4. 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.

in household income between the high-ranked and rising groups and the low-ranked and falling groups. The patterns we find are consistent with higher-earning individuals marrying or cohabitating 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.

6.4. Heterogeneity within groups

We find further evidence of segmented wealth mobility within clusters by examining the typical trajectories for each of the main subgroups that compose our baseline groups. These subgroups are an additional outcome from our hierarchical clustering and correspond to the first three branches of each group in the dendrogram presented in Figure 3B. We report the wealth rank trajectories of these subgroups in Figure 10. For convenience, we also reproduce the typical trajectory of each of the baseline groups in

light pink. The typical trajectories of each subgroup inform us of how diverse are the paths of individuals in the top, bottom, and middle of the wealth distribution.

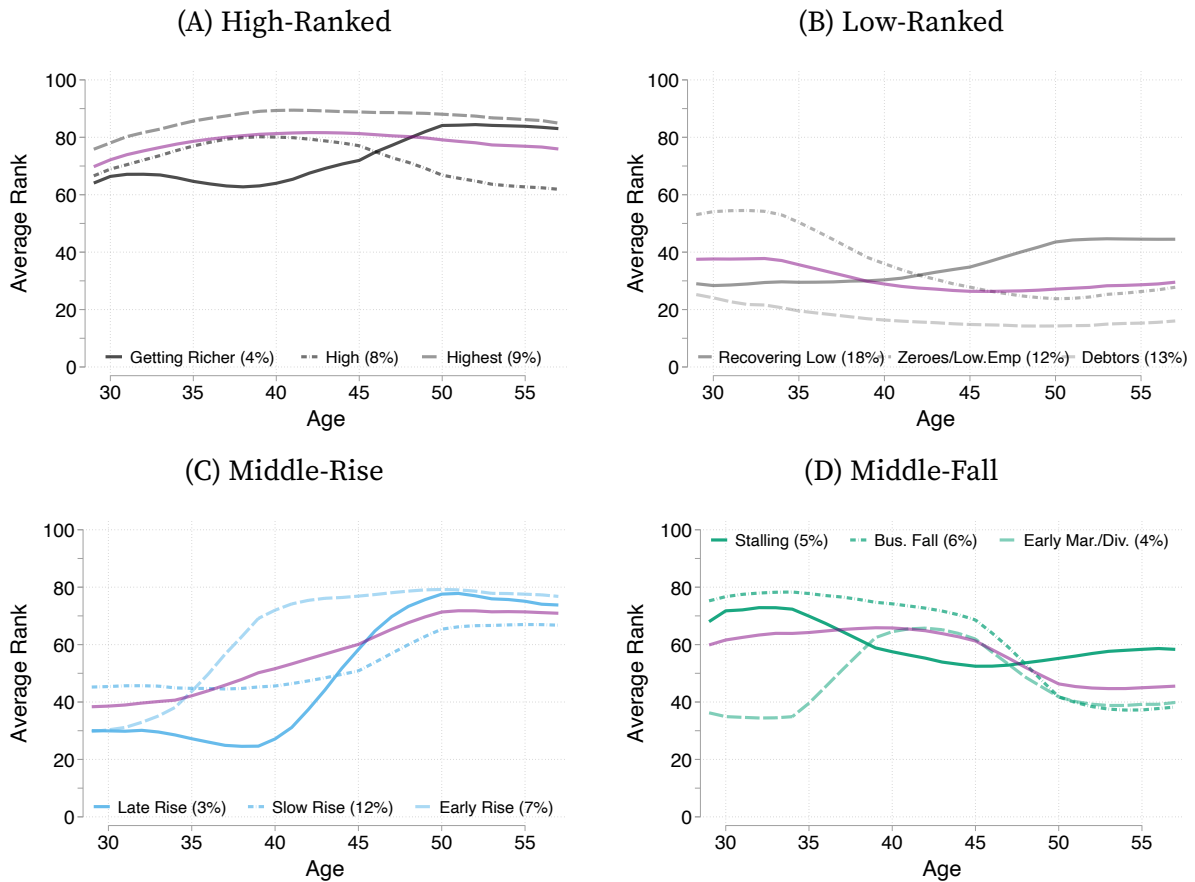
Overall, mobility within each main group remains contained within segments of the wealth distribution, even though there is overlap across groups. Specifically, mobility is limited to the upper part of the wealth distribution for the high-ranked group, the lower part for the low-ranked, and the middle for the risers and fallers. Thus, the trajectories of the subgroups reflect the dispersion of trajectories shown by the interquartile range in Figure 4. We now turn to describing the wealth mobility patterns that emerge for each subgroup. We examine their characteristics further in Appendix B.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.²⁹ 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 10A, have a larger fraction of their wealth in property (although homeownership rates are similar across the 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 distribution throughout their lives, and two groups that swap relative positions. We find

²⁹The 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 10. 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 4. 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 3B.

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 10C). We discuss later, in section 7, how

these timing differences are related to the education choices of the individuals of each subgroup. 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 cluster 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 aged. 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".³⁰ 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.

6.5. Summary of ex post group characteristics

The characteristics of the four main wealth clusters can be summarized as follows. The high-ranked individuals are homeowners early on and more likely to have business income. They also have the largest income levels across all groups. Conversely, the individuals who remain at the bottom of the wealth distribution have the lowest labor income and are less likely to own a home. Individuals falling through the wealth distribution are more likely to have business income and to be homeowners at age 30, but they have low household income. Finally, individuals rising through the wealth distribution have higher labor income relative to the fallers (particularly so at the household level) and become homeowners in their 30s and 40s.

³⁰The continuous decline of this group is accentuated by a substantial drop when its members are 45 to 50 years old. 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.

7. Parental background, education, and mobility

We next consider individuals' ex ante characteristics to determine whether the ex post heterogeneity in life-cycle trajectories across groups can be predicted by factors observed early in life. We specifically include characteristics emphasized in the intergenerational mobility literature, such as parental wealth and birthplace (Boserup, Kopczuk, and Kreiner 2017; Chetty et al. 2014) We also include education as another key determinant of labor income and intragenerational mobility.

7.1. Ex ante determinants of group membership

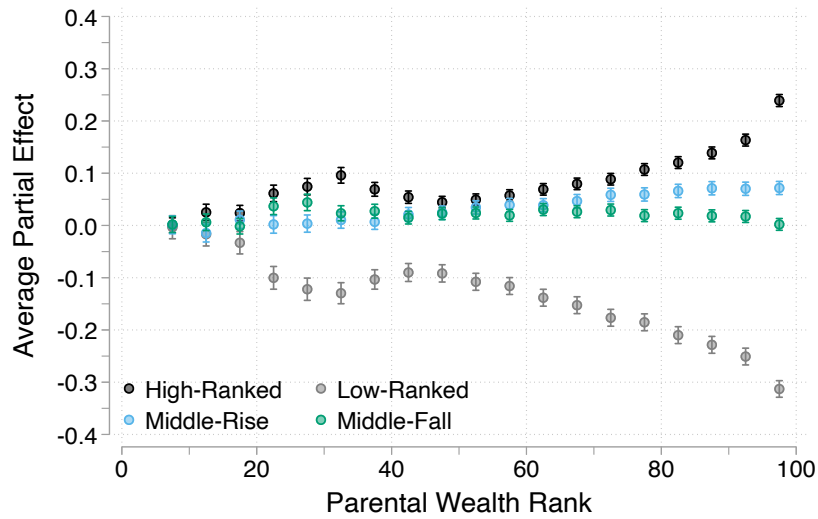
We quantify the predictive power of parental wealth, education, sex, and place of birth for group-assignment using a multinomial logit specification

$$\Pr (g_i = j) = F \left(\alpha_0^j + \beta_{q(i)}^j + \gamma_{\text{educ}(i)}^j + \delta_{\text{subj}(i)}^j + \lambda_{\text{male}(i)}^j + \mu_{\text{bcounty}(i)}^j \right), \quad (5)$$

where $F(\cdot)$ denotes the logit transformation. Specifically, we include ventiles of parental wealth fixed effects, $\beta_{q(i)}^j$.³¹ We also include education fixed effects (post-compulsory high school, technical college, undergraduate, postgraduate, 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. In practice, we aggregate the 350 degree-specific education codes into six categories: arts and humanities; business, economics, and agricultural management; computer science and engineering; natural sciences; health; and education specialists. Finally, we include a sex fixed effect, $\lambda_{\text{male}(i)}^j$, and birthplace fixed effects, $\mu_{\text{bcounty}(i)}^j$, that allow for place-based differences among the Oslo metropolitan area, other major cities, and rural regions, that are not captured by education choices and parental wealth.

³¹Formally, $q(i)$ 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.

FIGURE 11. Parental wealth rank and the probability of group assignment



Notes: The figure plots the average partial effect of Parental Wealth (measured in 1993) relative to being born to parents in the bottom ventile of the distribution. 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.

We find that parental wealth and education play a significant role in influencing group membership. However, we also find that the majority of the overall uncertainty over group membership at age 31 is not explained by these factors. This finding is to be expected because part of the value of our clustering approach, which exploits the full 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. We calculate the share of variation in group membership explained by (5) and decompose the partial contribution of parental wealth, education, and initial characteristics using a Shapley-Owen decomposition (Shorrocks 2013) in Appendix C.³²

Parental wealth. Figure 11 reports the average partial effects of parental wealth rank on predicted group assignment and their 95 percent confidence intervals. The role of

³²We describe the Shapley-Owen-Shorrocks decomposition in Appendix E. This approach allows us to calculate a single value per covariate category that is permutation-invariant and additively-decomposable.

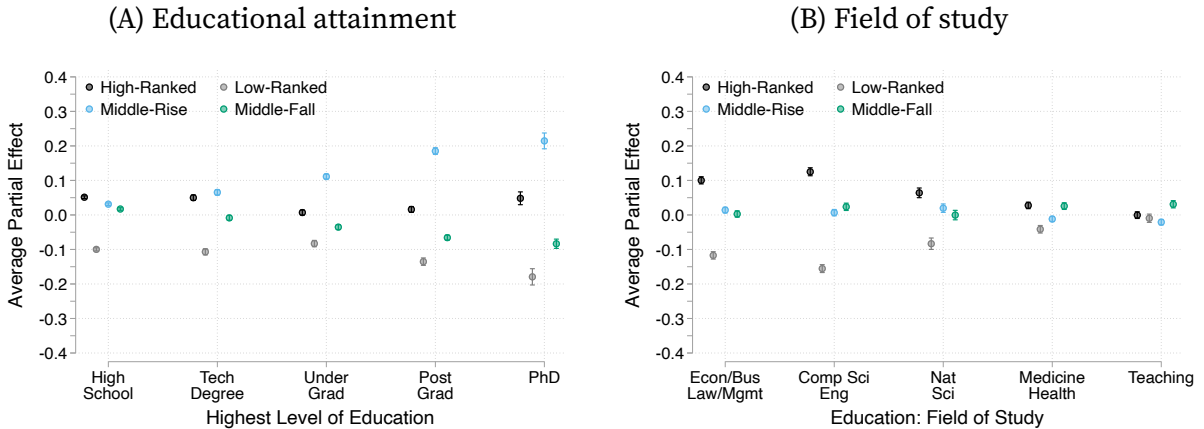
parental background is much stronger for the high- and low-ranked groups relative to the middle groups of risers and fallers. Individuals with wealthier parents are progressively more likely to belong to the high-ranked group and less likely to belong to the low-ranked group. An individual with parents in the top ventile of the wealth distribution is 25 percentage points more likely to belong to the high-ranked group than one with parents in the bottom ventile. The opposite is true for membership in the low-ranked group. By contrast, individuals coming from all sorts of 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.

Figure 11 shows that, on average, parental wealth has a large effect on group membership. Moreover, 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; see Table C.1 in Appendix C.3. 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.³³

Education. Figure 12 reports the estimated average partial effects of educational attainment and field of study, 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 high-ranked

³³In 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 (5) decrease 25 to 40 percent when conditioning on these extra variables, and the overall 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. We present these results in Appendix C.

FIGURE 12. Educational attainment and the probability of group assignment



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

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 up to 20 percentage points for PhDs, and provides evidence that risers are characterized by higher human capital accumulation that translates into increasing earnings profiles and wealth accumulation, as documented in Section 6.

Interestingly, the field of study is not as relevant for distinguishing risers and fallers, while 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. This finding adds to the fact that 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.

Education accounts for almost 40 percent of the fit of the model in classifying individuals across groups, so that together with parental wealth they make up for the vast majority of the explanatory power of initial characteristics; see Table C.1 in

Appendix C.3. Education variables are most relevant for the fit of the model in classifying risers, where it accounts for almost 55 percent of the fit; it accounts for almost 30 percent for the high-ranked group.

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. The birthplace of individuals has a small (but significant) effect on the probability of group membership. Being born in one of the four largest Norwegian cities increases the probability of being in the high-ranked or rising groups by close to 5 percentage points. Together, these variables explain less than 20 percent of the fit of the model in classifying individuals across groups, mostly because of sex differences in the groups' composition. We report detailed results in Appendix C.

7.2. Relationship to intergenerational mobility along the life cycle

Our results show that parental background plays an important role in shaping the life-cycle trajectories of individuals who are persistently at the top or bottom of the wealth distribution; it is comparatively much less informative about individuals who rise and fall through the middle of the distribution. This finding suggests that single-age snapshots of intergenerational wealth persistence might understate inequality of opportunity by not capturing the persistent effect of parental wealth on some individuals.

Moreover, the limited effect of education on the probability of belonging to the high-ranked group is noteworthy, especially when contrasted with the large effect of parental background. While parental background is key to explaining why an individual is persistently wealthy or persistently poor, education tells risers and fallers apart. We take this as evidence that there are substantial barriers to upward social mobility in

terms of wealth accumulation in Norway given that the high-ranked group makes up most of the top 20 percent of the individuals in this cohort and the group of risers remains below them throughout the sample (consistent with the pattern of segmented mobility in Figure 4 and 5).

These patterns result in a trend of declining intergenerational mobility along the life cycle that mirrors the increasing trend of intragenerational mobility. We formalize this in Appendix D, where we compute measures of rank persistence for the ranks of individuals in our sample with respect to their parents' wealth rank in 1993, analogous to the ones we introduced in Section 3. We find that the correlation between individuals' wealth ranks and their parents' wealth ranks increases with age, so that individuals become more similar in wealth to their parents as they age.

Interestingly, this trend of declining intergenerational mobility is common to all the main groups, and particularly strong for the risers. In fact, the relationship between parental- and own-ranks increases precisely when the reversal of fortunes of the risers and fallers takes place (right after age 35 when the risers' wealth rank starts to increase and after age 45 when the fallers' rank declines strongly). In this way, both the trend of declining intergenerational mobility and the trend of increasing intragenerational mobility are driven by the reversal of fortune of the middle groups.

Taken together, our results are consistent with the evidence on intergenerational wealth mobility measured by rank-rank persistence at fixed ages. Like [Boserup, Kopczuk, and Kreiner \(2017\)](#) we find increasing intergenerational persistence between ages 30 and 45, however, as we document above there is still considerable scope for wealth mobility over people's lifetimes. They are also consistent with the lifetime income profiles of more educated people being steeper but also exhibiting more variance in their slope (see, for example, [Guvonen 2009](#)). Finally, these results are in line with prior works showing that both wealthier and more educated individuals save more on average ([Dynan, Skinner, and Zeldes 2004](#); [Fagereng, Holm, Moll, and Natvik 2019](#)) as

well as results showing the importance of these margins in explaining lifetime inequality (Huggett, Ventura, and Yaron 2011).

8. 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 find similar qualitative patterns of segmented wealth mobility for alternative clustering algorithms and outcomes. In particular, we find similar patterns of mobility for the typical trajectories of wealth ranks recovered with these alternative methods. We view this as building credibility in our main results and the widespread applicability of the method we propose.

8.1. 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 of a lower dimensional group.³⁴ These tools are designed to summarize paths (or sequences) of categorical outcomes. Our approach is better suited to the continuous variation in wealth ranks we consider because (i) it avoids categorizing ranks into arbitrary discrete groups and (ii) it lets us exploit the cardinality of our rank

³⁴These 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).

measurements. Two individuals located at the 95th and 5th percentiles of the wealth distribution are further apart than two individuals located at the 95th and 50th percentile of the wealth distribution, for example. 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 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 . While the solution is typically fast for a given initial guess of the partition, the computational demands become substantial for a large number of observations.³⁵ 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 $G \in \{1, \dots, N\}$ 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 probabilistic assignments of individuals to groups. For instance, [Lewis, Melcangi, and Pilossoff \(2021\)](#) 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

³⁵Given the data security requirements for using datasets such as ours, there are limitations to the computational resources we can access.

workers to these types using the EM-algorithm.

8.2. 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 13, 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.³⁶ Similarly, our results are preserved when using the K -means algorithm to cluster individuals.

We begin by performing two robustness exercises where we group individuals based on their log-wealth histories and the trajectories of their “Lorenz” ordinates (which we define shortly) using our baseline clustering method.³⁷ The log-wealth is a concave transformation of the wealth levels and magnifies differences for individuals with relatively low levels of wealth, relative to those with high levels. The Lorenz ordinates, given by an individual’s position in the cohort’s Lorenz curve,³⁸ measures the distance between individuals by the share of total wealth that lies between them. Formally,

$$y_{i,t}^{Lorenz} = \frac{\sum_{\{j:w_{j,t} \leq w_{i,t}\}} w_{j,t}}{\sum_j w_{j,t}}. \quad (6)$$

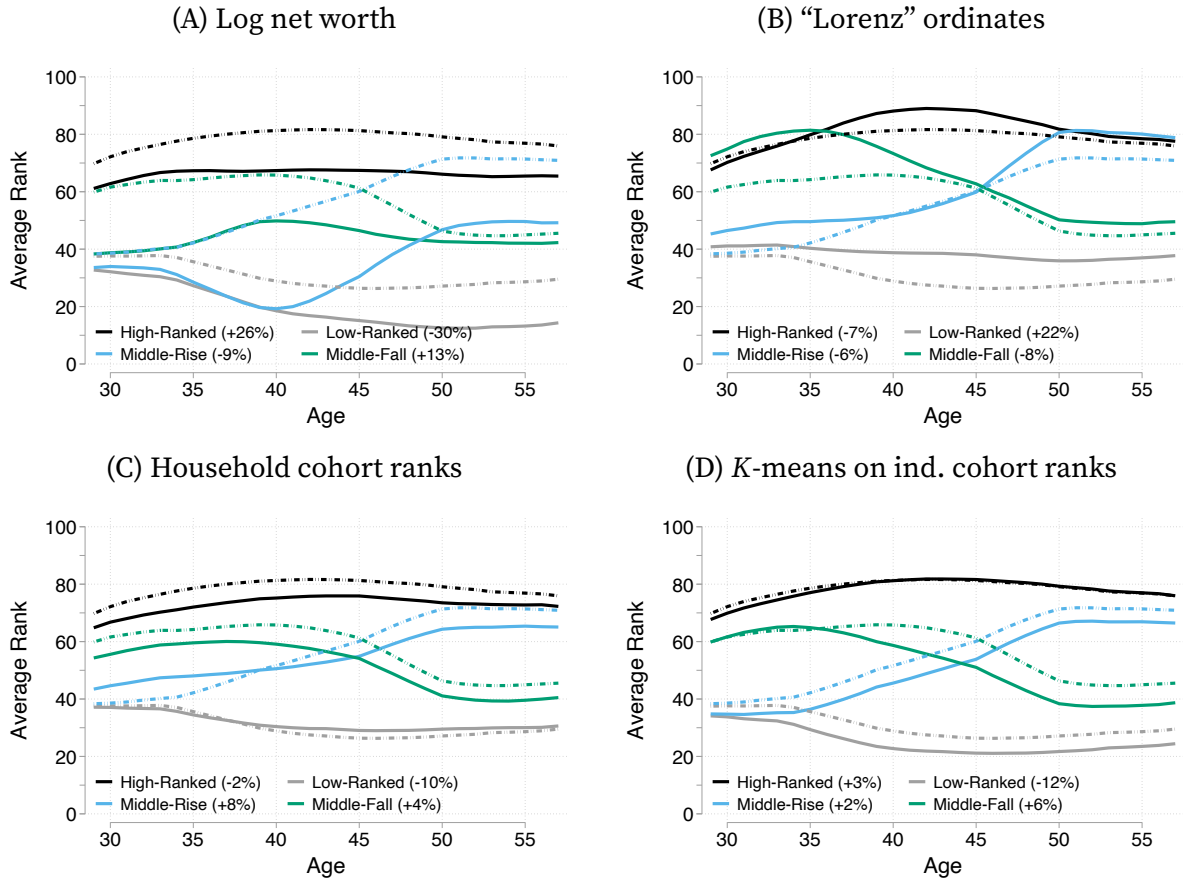
This is a convex transformation of wealth ranks and hence magnifies differences between individuals at the top of the distribution, making inequality more salient than either wealth ranks or log-wealth.

³⁶It is also not a mechanical result of clustering. 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.

³⁷In 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, making those individuals look more alike when clustering.

³⁸The Lorenz curve maps each (ranked) individual to the cumulative share of wealth held by individuals who are poorer than they are.

FIGURE 13. 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. 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). 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 4.

As Figures 13A and 13B show, these alternative clustering exercises do not change the defining mobility patterns described in Figure 4.³⁹ 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

³⁹All the panels of Figure 13 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.

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) and the remaining groups smaller with a change in composition toward higher-ranked individuals. Nevertheless, the patterns of the trajectories of each group are preserved.

In Figure 13C, 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 cohabitating individuals who make up about 80 percent of the sample throughout (see Figure 9A).

Finally, the exercise shown in Figure 13D 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, whose wealth rank profile also becomes on average lower.

9. 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. We found intragenerational mobility increases along 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 in the same birth cohort: 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 happen, 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. These relationships are, in turn, of interest for the validation of life-cycle models of wealth accumulation, in which these different mechanisms feature prominently.

Our results suggest several interesting directions for future work. We highlight two. First, like most economies, Norway is characterized by a large degree of wealth inequality. Yet, it is unclear whether the same wealth mobility patterns would arise under less generous welfare transfers, especially given our finding that many individuals accumulate very little wealth as they age. A second interesting and related avenue would be to relate wealth mobility to the incidence of wealth taxation over the life cycle.

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/eoj.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>.
- Becker, Gary S. and Nigel Tomes. 1979. "An equilibrium theory of the distribution of income and intergenerational mobility." *Journal of Political Economy* 87 (6):1153–1189. URL <http://www.jstor.org/stable/1833328>.
- 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, 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>.
- 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.
- . 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/eoj.12496>.
- 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 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](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>.
- 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 <https://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>.
- 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>.
- Guvonen, Fatih. 2009. “An empirical investigation of labor income processes.” *Review of Economic Dynamics* 12 (1):58–79. URL <https://doi.org/10.1016/j.red.2008.06.004>.
- Guvonen, 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>.
- Guvonen, 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.
- 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://johnrichumphries.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>.

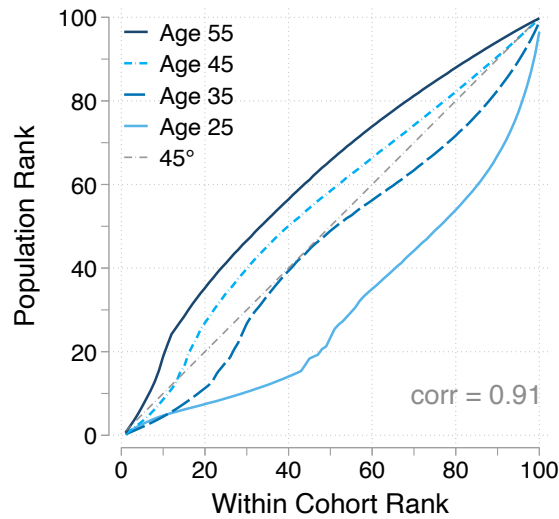
- Krueger, Alan B. 2012. “The rise and consequences of inequality in the United States.” URL https://obamawhitehouse.archives.gov/sites/default/files/krueger_cap_speech_final_remarks.pdf.
- Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2020. “Income and wealth inequality in America, 1949–2016.” *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.
- 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>.
- 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>.
- 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>.
- 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>.
- 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.

Appendix A. Additional results on measurement of wealth ranks

The dispersion of wealth within our cohort of interest increases as the cohort ages. This means that the correspondence between wealth and wealth levels changes over time. Figure A.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 A.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.

Appendix B. Additional results on typical trajectories

B.1. R² measures for partitions

We use three different measures of the R² to determine the share of variation in wealth rank trajectories explained by the clustering for a given partition into $G > 1$ groups. Our measures, which decompose variation in ranks into variation within an individual's trajectory and variation between individuals' trajectories, are

$$R^2_{overall} = 1 - \frac{\sum_i \sum_t (y_{i,t} - \bar{y}_t^{g(i)})^2}{\sum_i \sum_t (y_{i,t} - \bar{y})^2} \quad \text{overall } R^2 \quad (\text{B.1})$$

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

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

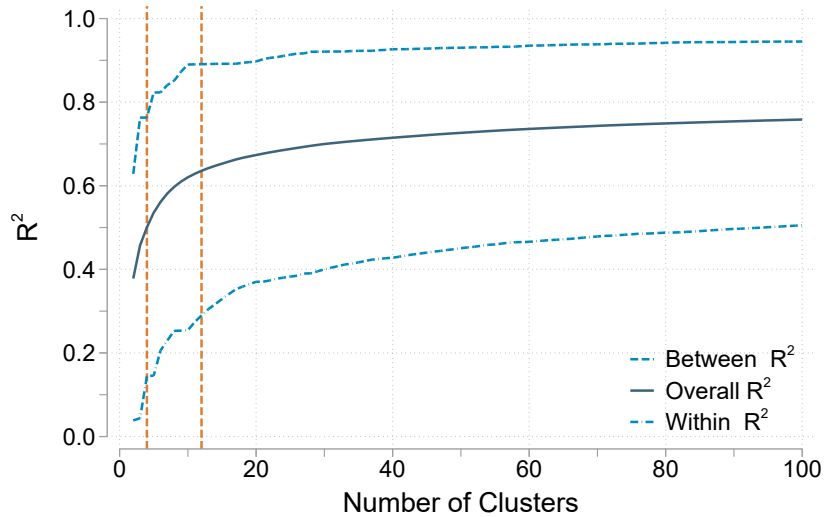
The overall R² measures the share of variation in ranks relative to their unconditional average over the longitudinal and cross-sectional dimensions, \bar{y} , explained by the clusters' typical trajectories, $\bar{y}_t^{g(i)}$. We denote an individual's assigned cluster by $g(i)$.

The between R² 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{y}_i its average across individuals, and $\bar{y}_i^{g(i)}$ its average for cluster $g(i)$.

The within R² 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 $\bar{y} - \bar{y}_i$. This gives a measure of the total *within* variation. The cluster's deviation, $\bar{y}_t^{g(i)} - \bar{y}^{g(i)}$, where $\bar{y}_t^{g(i)}$ is the cross-sectional average of ranks for cluster $g(i)$ in time t and $\bar{y}^{g(i)}$ is its average over time is used to construct a measure of the explained within variation produced by the clusters' typical trajectories.

Figure B.1 presents the three R² 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 R² even when clusters increase. In contrast to this, the between R² 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 B.1. R^2 Measures, up to 100 clusters



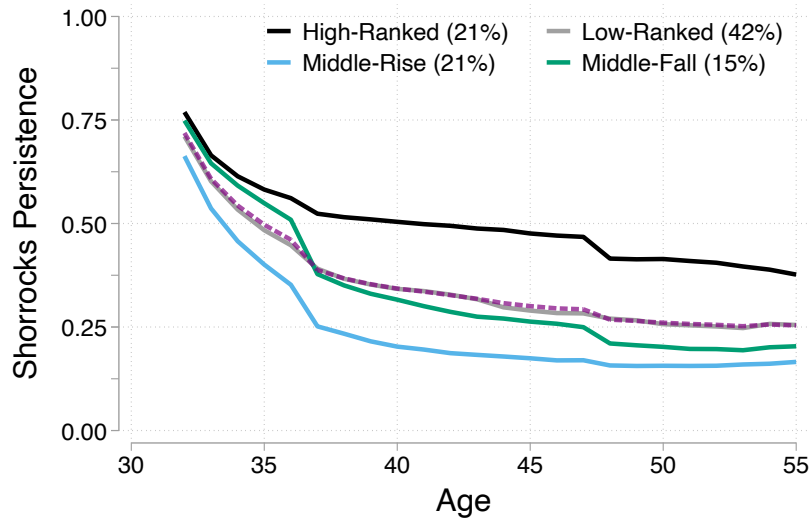
Notes: The figure plots the share in the variation of wealth rank trajectories explained by the partitions induced by agglomerative hierarchical clustering algorithm for $G = 1, \dots, 100$ groups. The overall R^2 , defined in equation (B.1), and also presented in Figure 3A. The between R^2 , defined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

B.2. Intragenerational mobility by group: Shorrocks

Analogously to Figure 6, we also decompose the Shorrocks mobility measure reported in Figure 2 by group. To do so, we first define population-level quintiles identically to Figure 2; doing this allows us to compare mobility across clusters. We then 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. We report these results in Figure B.2. 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, reflecting large changes in ranks over the life cycle, with risers experiencing the highest mobility.

FIGURE B.2. 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.

B.3. Further group characteristics

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

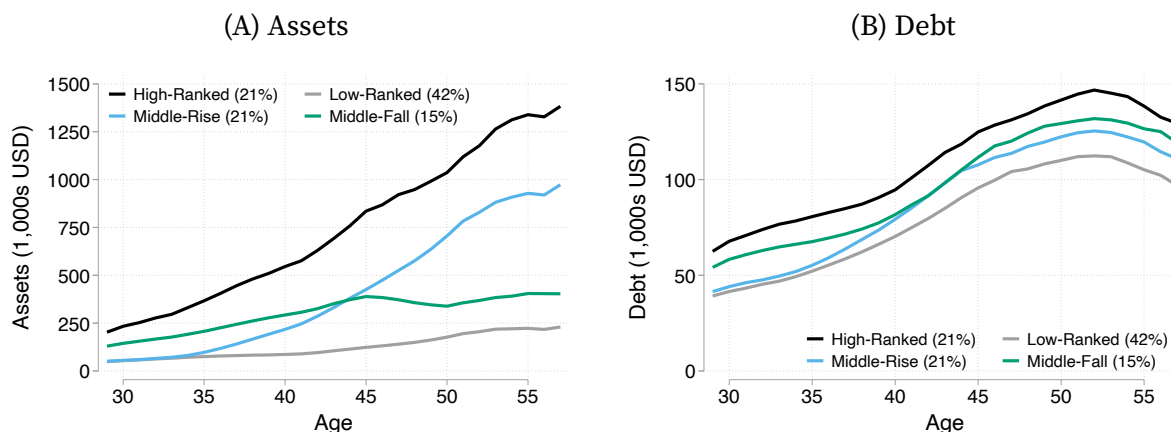
Figure B.3 complements Figure 7 by providing life-cycle profiles for the asset and debt levels of the four groups. The main takeaway is that the differences in net-worth reported in Figure 5 mostly reflect differences in assets across groups. The magnitude of debt and dispersion of debt are much lower than those of assets.

The business ownership rates are much more stable throughout the sample and show a clear ordering of the groups. The high-ranked have a markedly higher business ownership rate, followed by that of the fallers. The risers' business ownership rate increases somewhat but remains below the first two groups. The low-ranked group has the lowest business ownership rate.

Figures B.4 and B.5 complement Figure 8 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 B.5A shows that the differences in employment income among groups follow

FIGURE B.3. Asset and debt trajectories by group

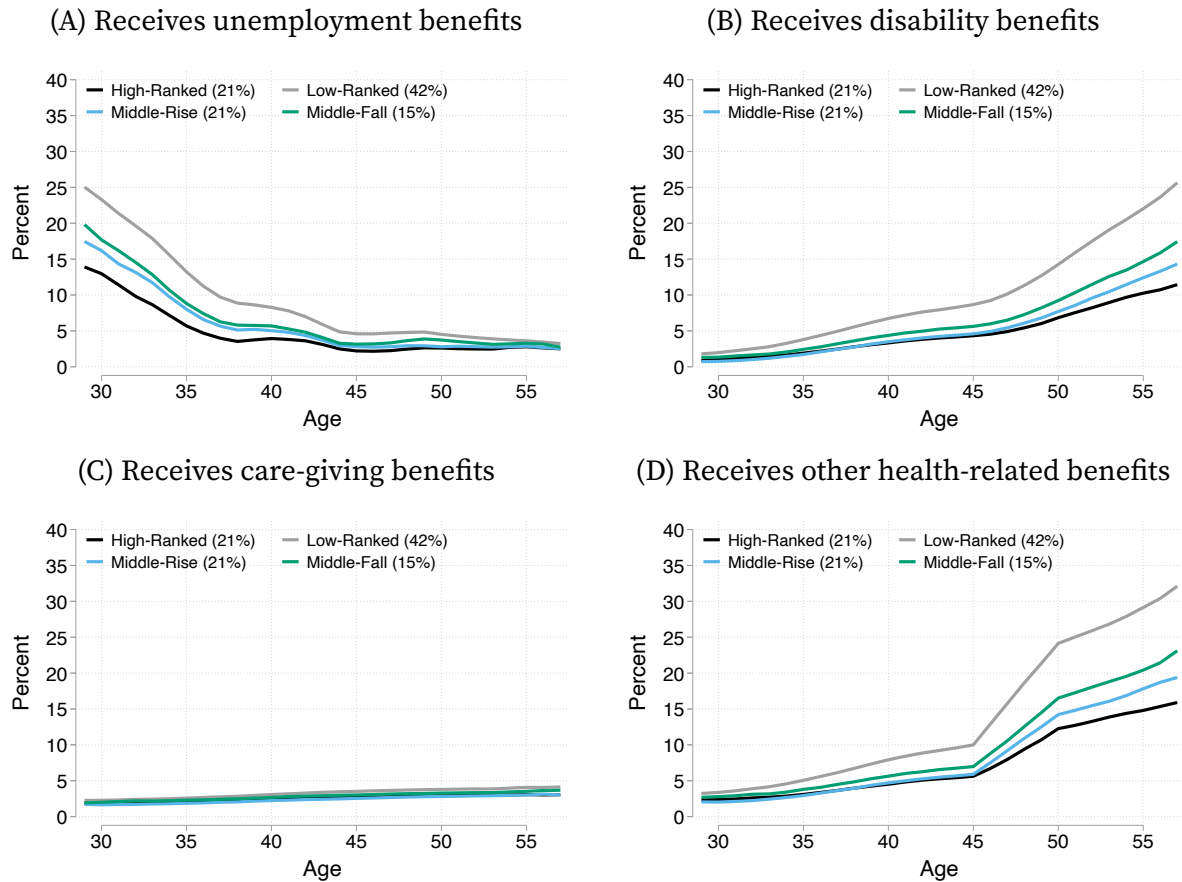


Notes: Panels A and B plot, respectively, average asset and debt trajectories in 2019 US dollars of the four main groups presented in Figure 4.

the same trends as the differences in overall income shown in Figure 8B but are not as large. The discrepancy lies in the differences in self-employment and capital income. Figure B.5B 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 8D). 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).

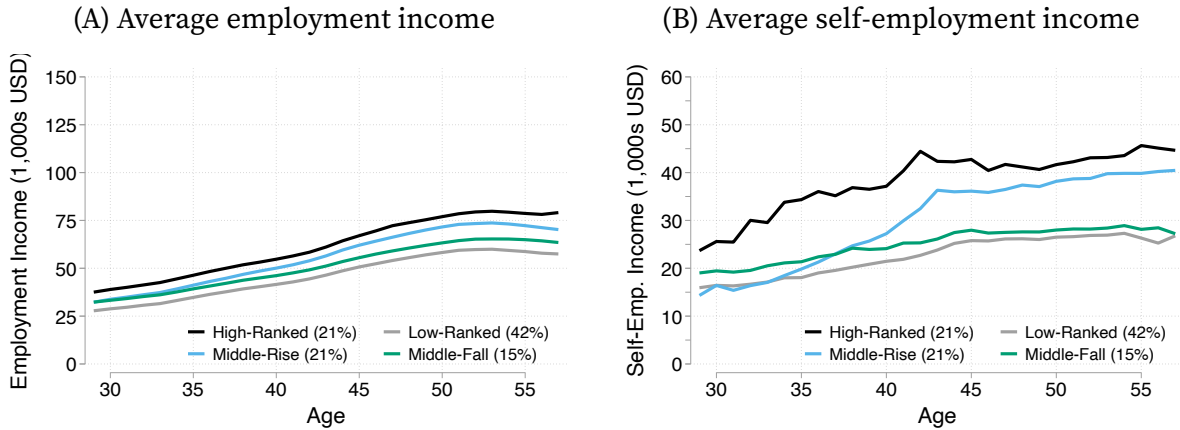
Finally, Figure B.6 presents the life cycle profile of household wealth ranks for the individuals in each of the four main groups of the 1960–64 birth cohort. Household wealth ranks have the same qualitative (and even quantitative) behavior as individual wealth ranks. To construct households, we match 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.

FIGURE B.4. Benefit receipt by group



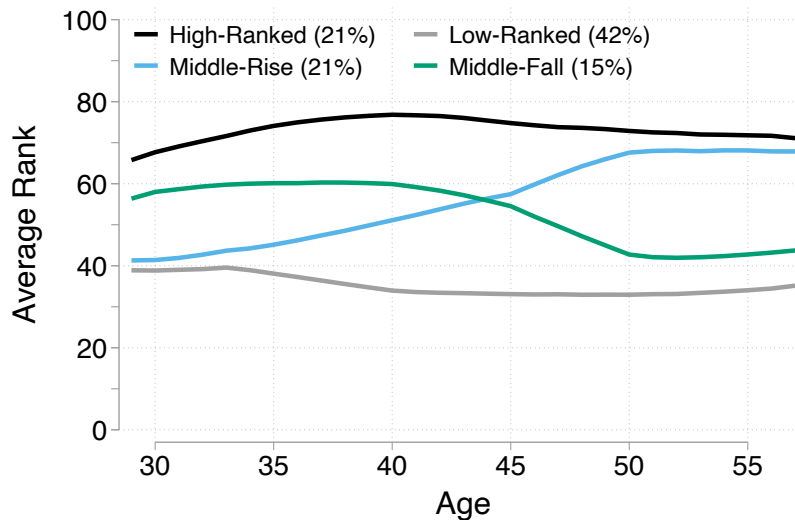
Notes: The figures plot the share of individuals by age that receive benefits from public programs for the four main groups presented in Figure 4. We consider four types of benefits: unemployment; disability; transfers related to care-giving including those for parents of young children, elderly care, 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 B.5. 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 4. 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 B.6. 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 4. 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.

B.4. Group characteristics for subgroups

We provide here additional results characterizing the subgroups of our four main groups, described in Section 6.4. We describe the main patterns below.

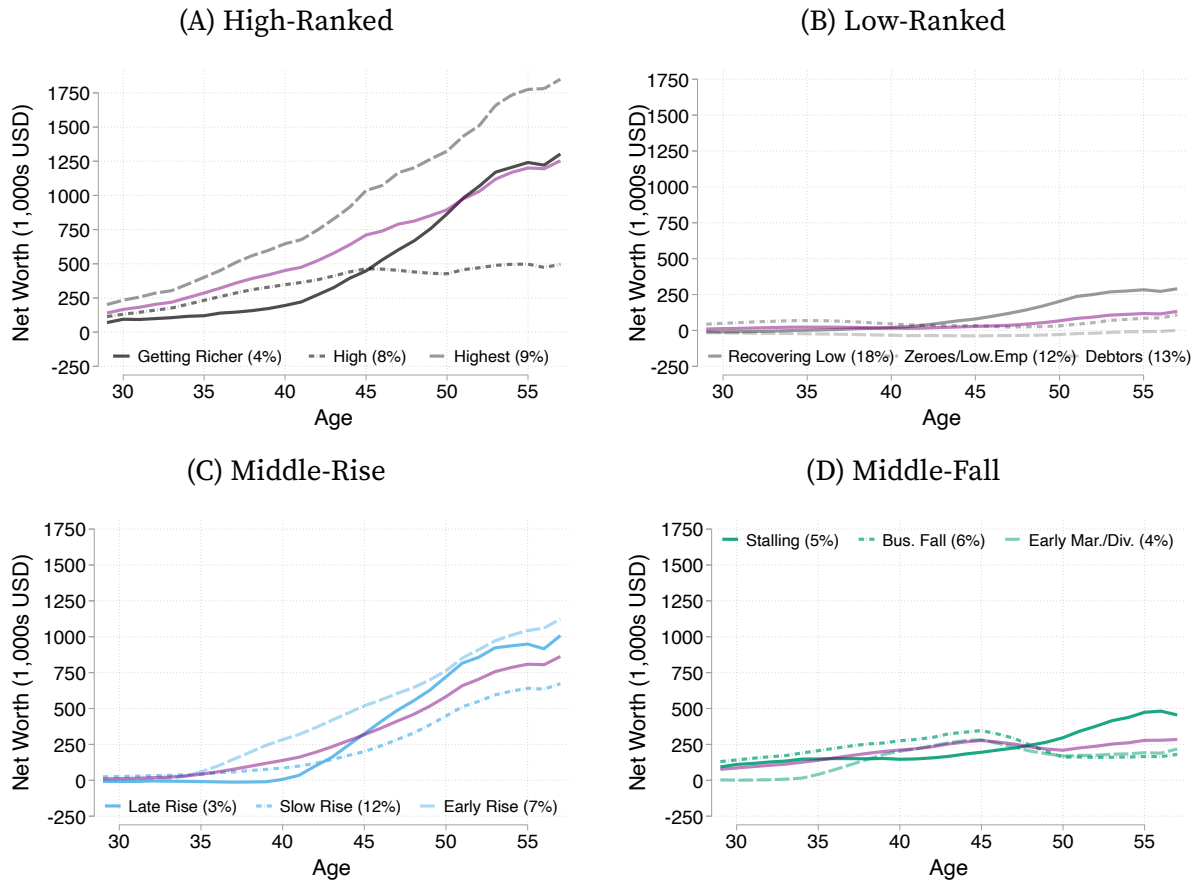
The typical mobility patterns of these groups are presented in Figure 10 in the main text. The corresponding typical wealth trajectories are in Figure B.7. The common scale in the figure highlights the vast differences in net-worth across subgroups. These are much larger than the differences across groups 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 B.9). The differences among the fallers are less pronounced. The early marriage/divorce group stands out because of the accumulation of property (again, Figure B.9). This group’s moniker follows from its household formation and dissolution dynamics.

Figures B.8 and B.10 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 B.11 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.

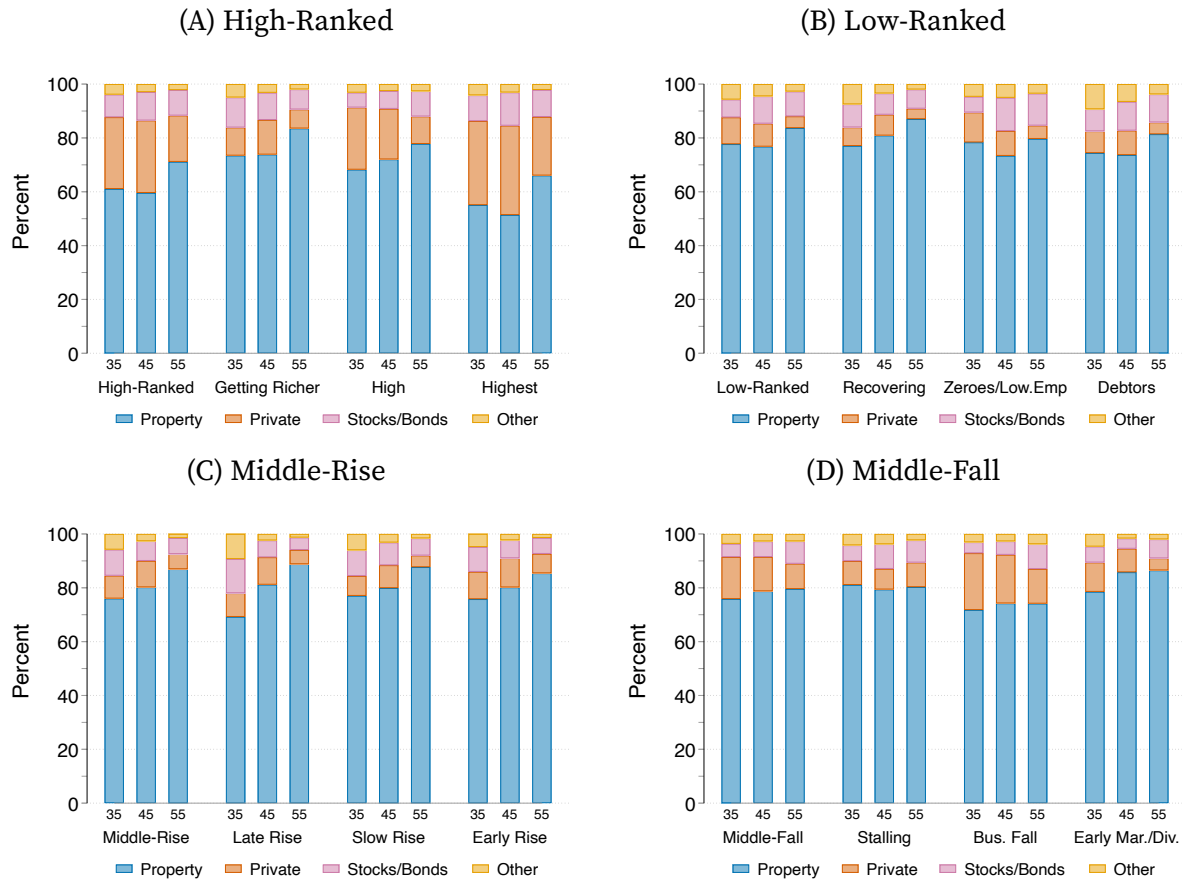
Finally, Figure B.12 reports educational attainment for all four main groups and their subgroups. Risers (and their subgroups) have the highest educational attainment, even when compared with the high-ranked subgroups. Tellingly, the subgroup with the next highest education is the “getting richer” subgroup 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 B.7. Wealth levels by subgroup



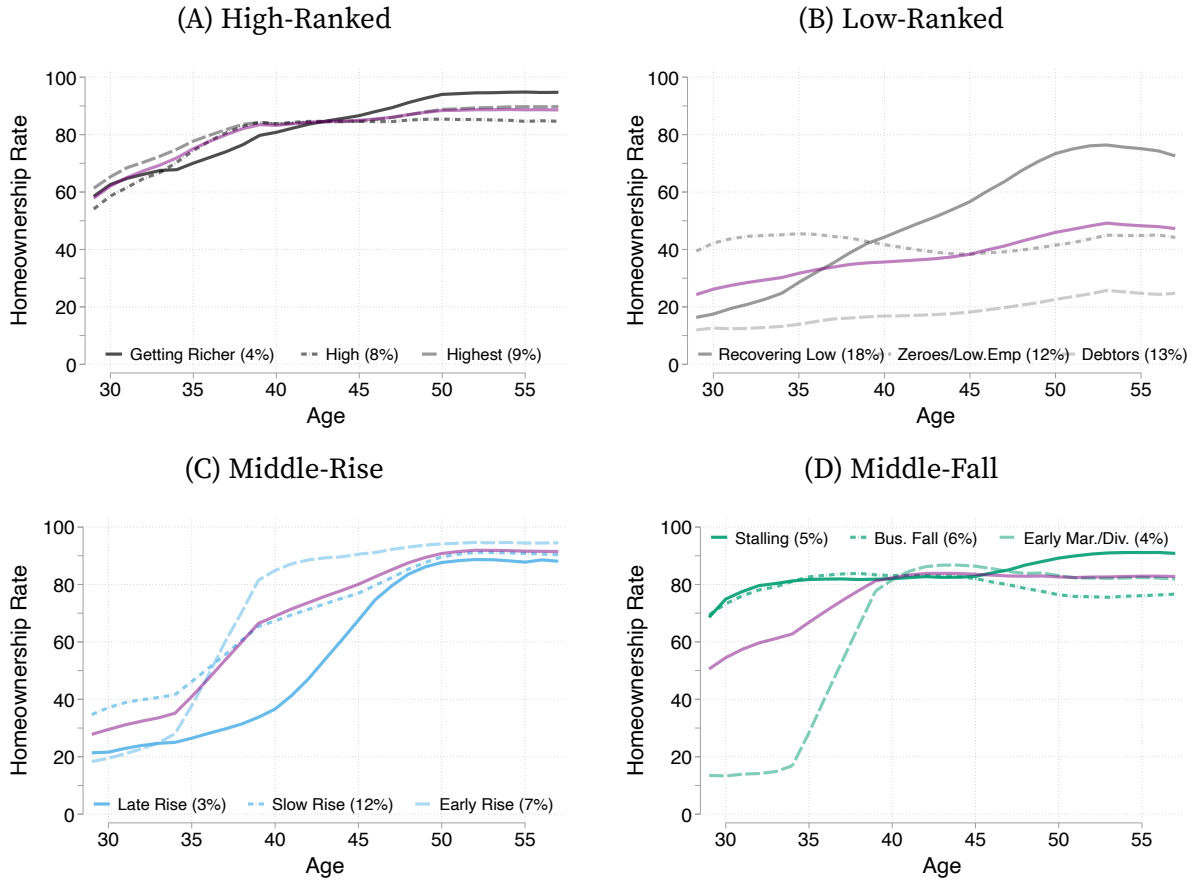
Notes: The figure plots the average wealth level in thousands of US dollars in each clustered subgroup against the cohort's average age. All individuals belong to the 1960–64 birth cohort. For each subgroup we additionally report the major group's values. This corresponds to Figure 5 in the main text.

FIGURE B.8. Portfolio composition by subgroup



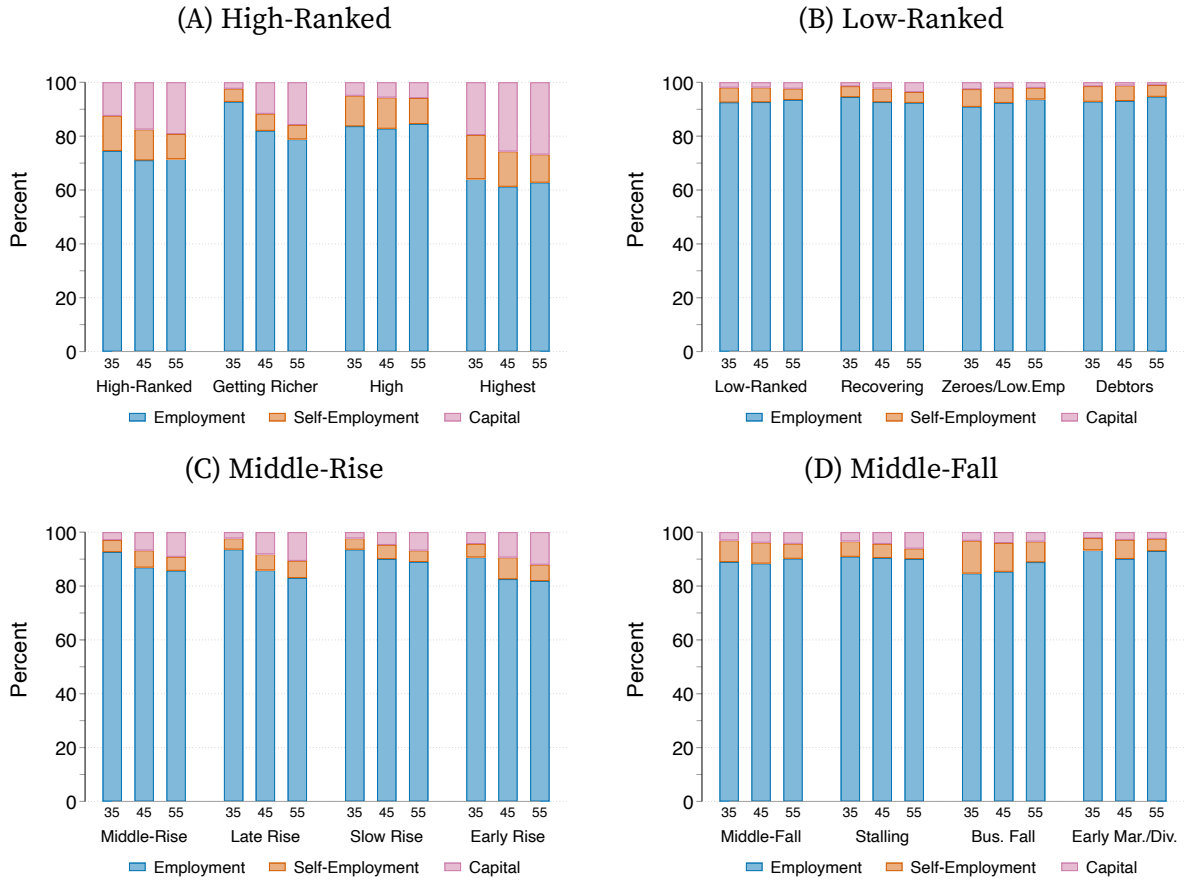
Notes: This figure reports the share of assets accounted for by property, privately held assets, financial assets, and other assets, defined as the total value of each asset class divided by the total assets within a group. For each subgroup, we additionally report the major group's values. This corresponds to Figure 7A in the main text.

FIGURE B.9. Homeownership rates by subgroup



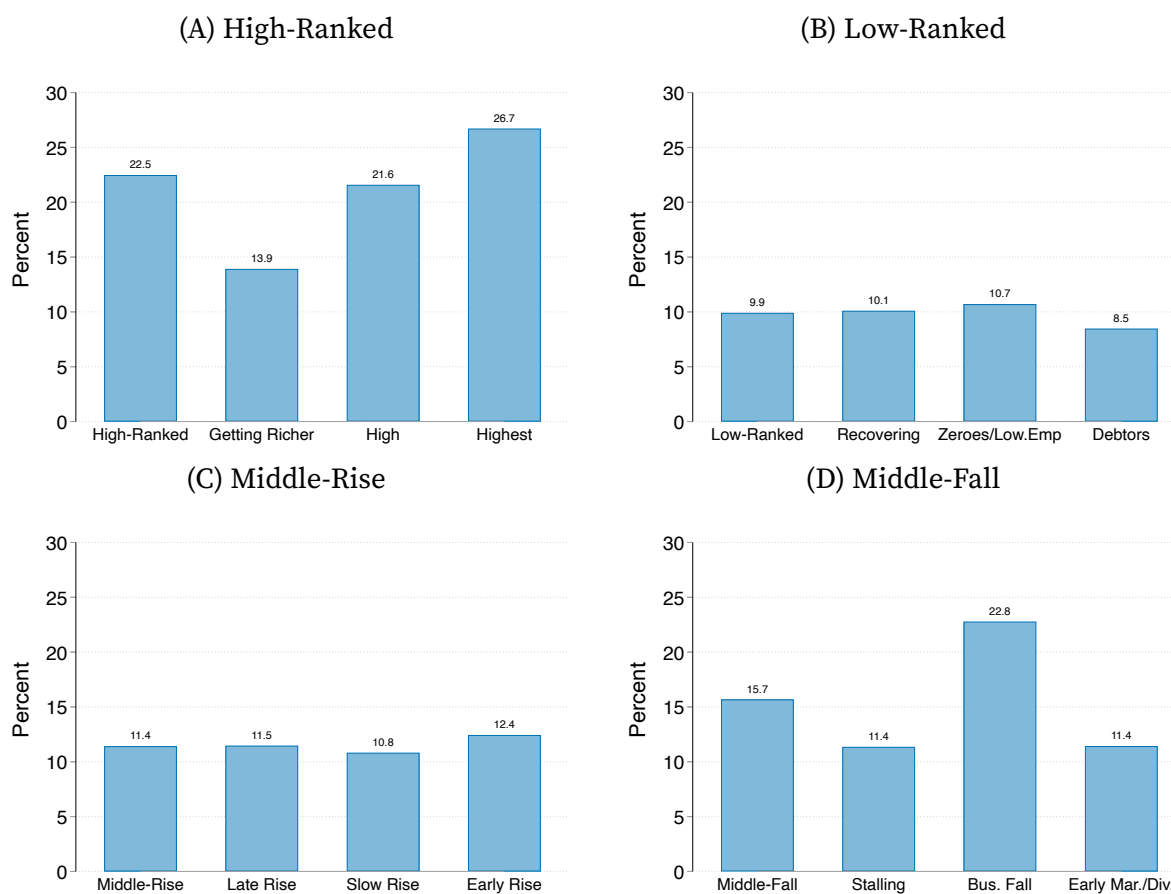
Notes: The figure plots the share of individuals who are homeowner in each clustered subgroup against the cohort's average age. For each subgroup we additionally report the major group's values. This corresponds to Figure 7B.

FIGURE B.10. Income composition by subgroup



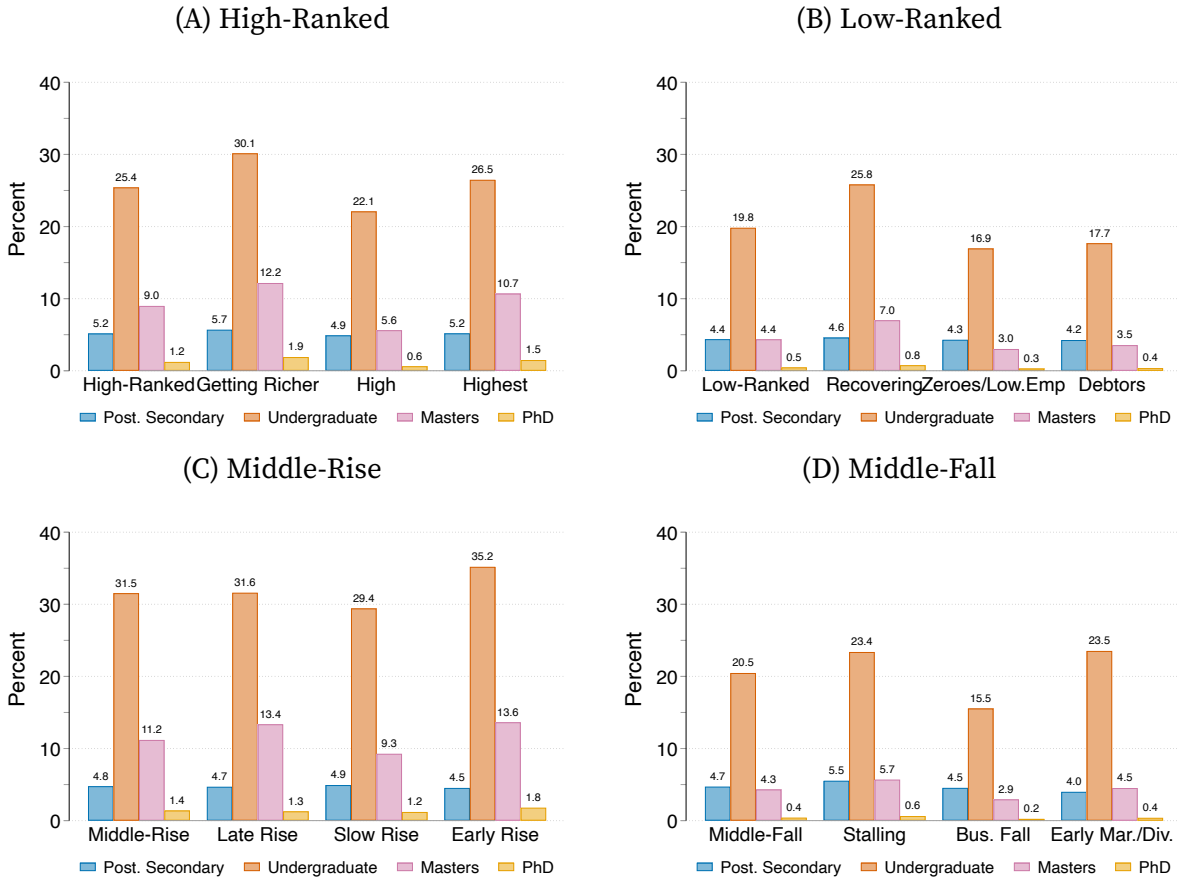
Notes: This plots the share of each group's income accounted for by employee, self-employment, and capital income. For each subgroup, we additionally report the major group's values. This corresponds to Figure 8C.

FIGURE B.11. 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 8D.

FIGURE B.12. 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 7.

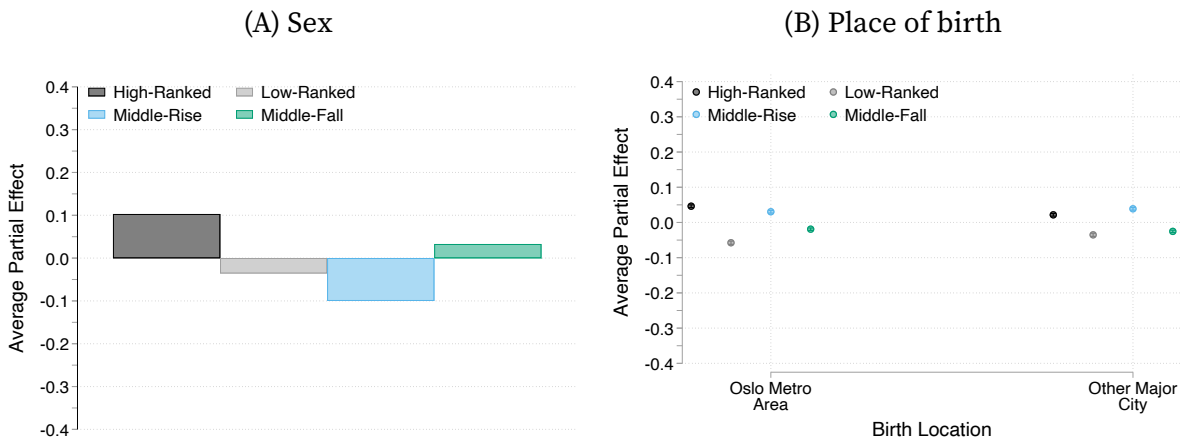
Appendix C. Additional results on ex ante analysis

C.1. The role of sex and birthplace

Sex. Panel A of figure C.1 shows the average partial effects associated with gender and birth place. 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 B of Figure C.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.

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

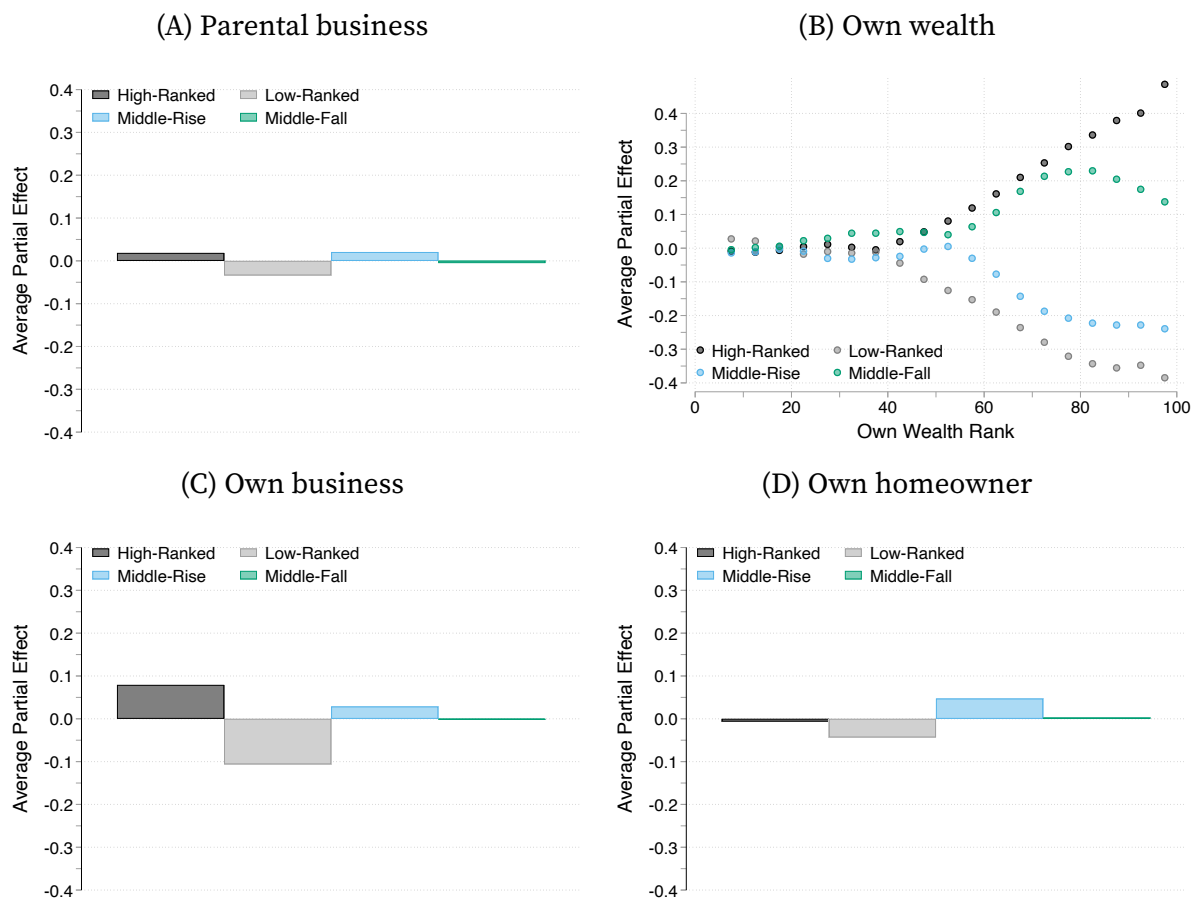


Notes: The figures plot the average partial effect of men relative to women (panel A) and urban areas relative to rural areas (panel B). We construct the average partial effect by integrating over the empirical joint distribution of other covariates. We report point estimates separately for each outcome, the probability of being assigned to each of our four groups, along with their 95 percent confidence intervals.

C.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

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

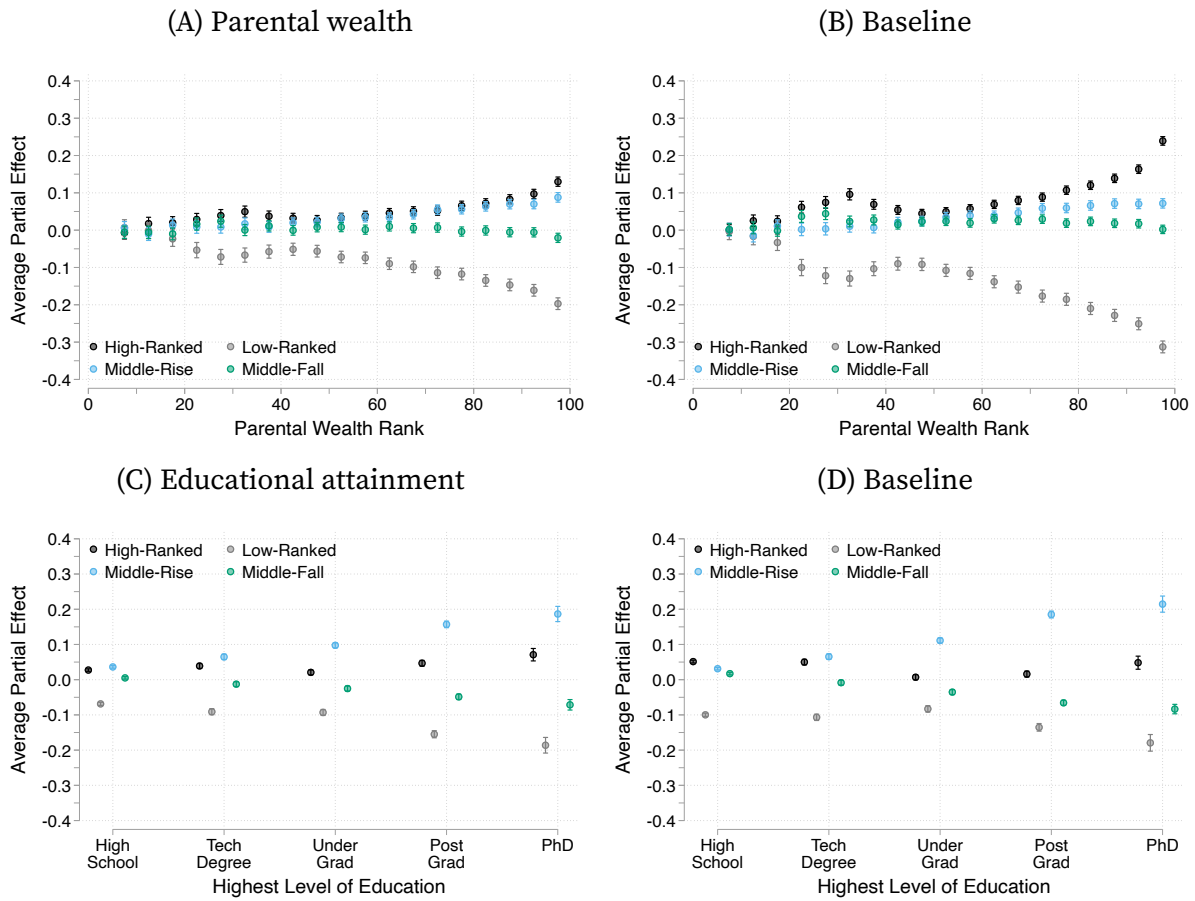


Notes: The figures plot the average partial effect on the probability of belonging to each of the four main groups of whether parents have a business in 1993 (panel A), the individual’s own wealth ventile in 1993 (panel B), whether the individual own a business in 1993 (panel C), and whether the individual 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 (5) 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.

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 entirely consistent with the patterns of segmented mobility we document in Section 5. We report the average partial effects for these additional controls in Figure C.2. The results

FIGURE C.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.

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 Table C.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 7, 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 C.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 remain significant and display the same qualitative patterns—confirming the relative importance of education and parental wealth among risers and the high-ranked.

C.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 (5) helps to explain the variation across groups.

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

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

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 (C.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

$$\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^G \widehat{\Pr}(g = k | X_i) \mathbb{1}[g(i) = k]. \quad (\text{C.2})$$

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 C.1 reports the total contribution from our four groups of ex ante regressors to

TABLE C.1. Predictive power of ex ante characteristics

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

Group	All Effects	Partial Contribution			
		Parent	Education	Sex	Birth Place
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	Birth Place
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 C.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 (5). 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 E.

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 E.

Consistent with the average partial effects we report in Figures 11 and 12, 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.

Relative to the results we report in the main text, Table C.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 (5) 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.

C.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 5) with the additional covariates described in Appendix C.2.

We present the results in Table C.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 C.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 C.2. Predictive power of ex ante characteristics with additional covariates

Total Contribution	Partial Contribution					
	Parent	Education	Sex	Birth Place	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 C.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 (5) with the additional covariates introduced in Appendix C.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 E.

Appendix D. Declining intergenerational mobility

An additional important aspect of mobility, that has been frequently studied in the literature, in particular, for income (see [Deutscher and Mazumder 2023](#) for a recent survey of different methodologies), is the extent to which the position in the distribution of children reflects that of their parents. This measure is informative about the intergenerational transmission of inequality and about equality of opportunity more broadly.

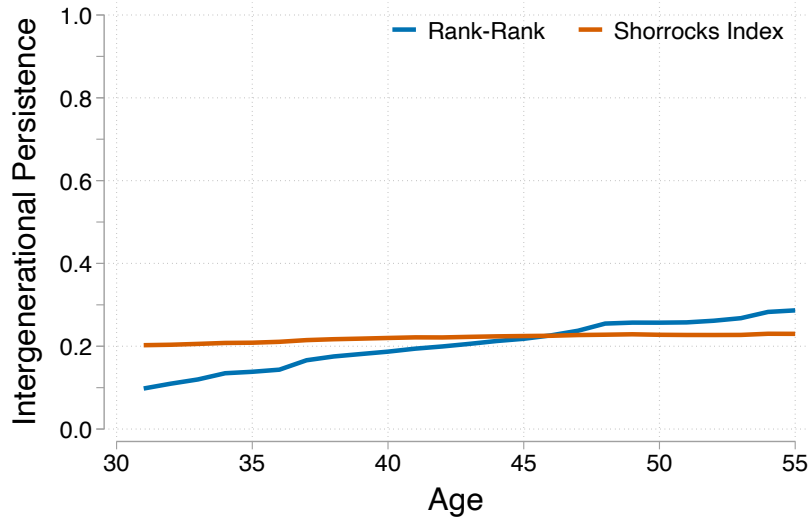
Our data also allows us to study mobility from an intergenerational perspective, which we show in [Figure D.1](#). This Figure represents the same persistence measures as [Figure 2](#), but regressing the current rank of our individual of interest on the rank of their parents when the individual was 30 years old. On average, this implies that we study persistence of wealth with respect to the wealth of the parents when they were 55. For the rank correlation measure, we regress $y_{i,t} = \alpha_t + \rho_t^p y_{p(i),1993} + u_{i,t}$, where $y_{p(i),1993}$ is the maximum of individual i 's parents' within-cohort wealth rank in 1993. For the Shorrocks index, we compute the share of individuals who are in the same wealth quintile as their parents were in 1993.

Interestingly, we find that the patterns of intergenerational persistence mirror those of intragenerational persistence. Intergenerational mobility decreases with age, captured by an increasing correlation between an individual's rank and their parents' rank. Put another way, we find that individuals become more similar in wealth to how their parents were as they age. This suggests that parental background conditions wealth trajectories beyond the level of initial wealth, as we explored in [Section 7](#). Interestingly, this pattern, as well as the magnitude of the parent-child rank-rank correlation coefficient, are very similar to the ones reported by [Boserup, Kopczuk, and Kreiner \(2017\)](#) in their study of intergenerational wealth mobility in Denmark.

Looking at the patterns of intergenerational mobility across groups ([Figure D.2](#)), we observe smaller differences than in the case of intragenerational persistence ([Figures 6 and B.2](#)). For all groups, persistence rises over the life cycle; individuals in every group becomes more similar to their parents in terms of wealth rank as they age. The level differences are mostly parallel, with the exception of risers.

Notably, the wealth rank of risers is negatively correlated with that of their parents at labor market entry. This correlation becomes positive over the life cycle. Thus, the members of the risers start out looking less like their parents (relatively to other groups), but they also converge faster to the wealth ranks of their parents over their lives.

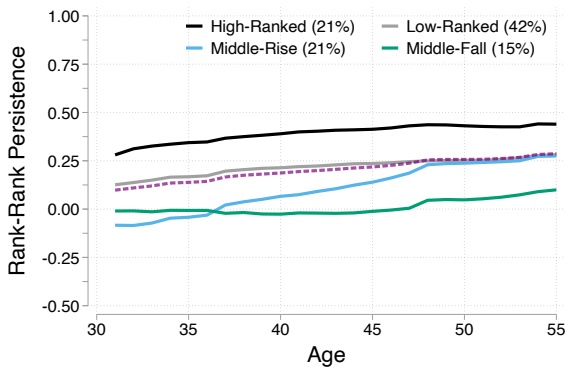
FIGURE D.1. Intergenerational persistence of wealth ranks



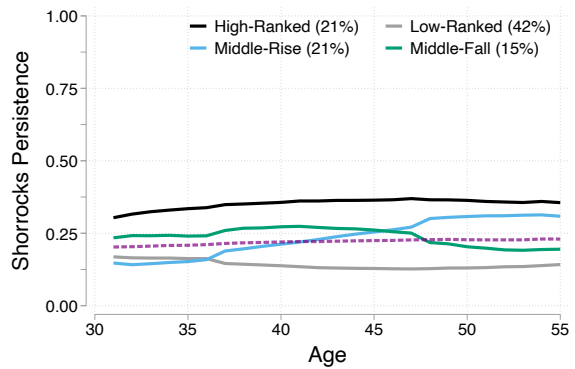
Notes: The figure shows the rank-rank and the Shorrocks' (1978) persistence measures for intergenerational mobility, using the level of wealth of parents when children were age 30 as a reference.

FIGURE D.2. Intergenerational persistence across groups

(A) Rank-Rank intergenerational Persistence



(B) Shorrocks intergenerational Index



Notes: Panel A plots the intergenerational rank-rank persistence for our four main groups. Panel B plots the intergenerational Shorrocks index for our four main groups. The pooled cohort-level persistence measures are shown in dashed lines.

Appendix E. The Shapley-Owen-Shorrocks Decomposition

Given an arbitrary function $Y = f(X_1, X_2, \dots, X_n)$, the Shapley-Owen-Shorrocks decomposition is a method to decompose the value of $f(\cdot)$ into each of its arguments X_1, X_2, \dots, 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 $f(\cdot)$,

$$\sum_{j=1}^n C_j = f(X_1, X_2, \dots, X_n), \quad (\text{E.1})$$

so that $C_j/f(\cdot)$ can be interpreted as the proportion of $f(\cdot)$ that can be attributed to X_j .⁴⁰ 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

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

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 .⁴¹ For example,

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

⁴⁰The 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 $C_j < 0$.

⁴¹We 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 $f(X_1, \emptyset_2, \dots, \emptyset_n)$, 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 decomposition in (E.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: $\underbrace{\{(0, 0)\}}_{k=0}, \underbrace{\{(1, 0), (0, 1)\}}_{k=1}, \underbrace{\{(1, 1)\}}_{k=2}$.

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

Nonlinear example

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

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

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 :

$$C_1 = \sum_{k=0}^2 \frac{(3-k-1)!k!}{3!} \left(\sum_{s \subseteq \mathcal{S}_k \setminus \{X_3\}; |s|=k} [f(s \cup X_1) - f(s)] \right) = \beta_1 X_1 \quad (\text{E.4})$$

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

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

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.

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

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,

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

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 $i = 1, \dots, M$ observations,

$$y_i = \mathbf{x}'_i \beta + u_i = \beta_0 + \sum_{j=1}^n \beta_j x_{ij} + u_i, \quad (\text{E.7})$$

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

$$\hat{y}_i = \mathbf{x}'_i \hat{\beta} = \hat{\beta}_0 + \sum_{j=1}^n \hat{\beta}_j x_{ij}, \quad (\text{E.8})$$

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

The function of interest is $f(X_1, \dots, X_K) = R^2$, defined as the explained sum of squares SSE over the total sum of squares SST

$$R^2(X_1, X_2, \dots, X_n) = \frac{SSE}{SST} = \frac{\sum_{i=1}^M (\hat{y}_i - \bar{y})^2}{\sum_{i=1}^M (y_i - \bar{y})^2}. \quad (\text{E.9})$$

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

$$R^2(\emptyset) = \frac{\sum_i^M (\hat{\beta}_0 - \bar{y})^2}{\sum_i^M (y_i - \bar{y})^2} = 0, \quad (\text{E.10})$$

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:

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

$$R_2^2 = \frac{1}{3} \left[R^2(X_2) - R^2(\emptyset) \right] + \frac{1}{6} \left(\left[R^2(X_1, X_2) - R^2(X_1) \right] + \left[R^2(X_2, X_3) - R^2(X_3) \right] \right) + \frac{1}{3} \left[R^2(X_1, X_2, X_3) - R^2(X_1, X_3) \right]; \quad (\text{E.12})$$

$$R_3^2 = \frac{1}{3} \left[R^2(X_3) - R^2(\emptyset) \right] + \frac{1}{6} \left(\left[R^2(X_3, X_2) - R^2(X_2) \right] + \left[R^2(X_1, X_3) - R^2(X_1) \right] \right) + \frac{1}{3} \left[R^2(X_1, X_2, X_3) - R^2(X_2, X_1) \right]. \quad (\text{E.13})$$

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

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

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.