American Economic Journal: Applied Economics 2024, 16(2): 1–30
https://doi.org/10.1257/app.20210137

Measuring Absolute Income Mobility: Lessons from North America and Europe†

By Robert Manduca, Maximilian Hell, Adrian Adermon, Jo Blanden, Espen Bratberg, Anne C. Gielen, Hans van Kippersluis, Keunbok Lee, Stephen Machin, Martin D. Munk, Martin Nybom, Yuri Ostrovsky, Sumaiya Rahman, and Outi Sirniö*

We use linked parent–child administrative data for five countries in North America and Europe, as well as detailed survey data for two more, to investigate methodological challenges in the estimation of absolute income mobility. We show that the commonly used "copula and marginals" approximation methods perform well across countries in our sample, and the greatest challenges to their accuracy stem not from assumptions about relative mobility rates over time but from the use of nonrepresentative marginal income distributions. We also provide a multicountry analysis of sensitivity to specification decisions related to age of income measurement, income concept, family structure, and price index. (*JEL* D31, G51, I31, J12, J31, J62)

The hope that standards of living rise from one generation to the next is widely shared across the world. In the United States, this goal is often considered part of the "American Dream." Over the past decade, amid growing concern that upward mobility has stalled in many high-income countries, there has been

* Manduca: Department of Sociology, University of Michigan (email: rmanduca@umich.edu); Hell: Department of Sociology, Stanford University (email: hell@stanford.edu); Adermon: Institute for Evaluation of Labor Market and Education Policy (email: adrian.adermon@ifau.uu.se); Blanden: School of Economics, University of Surrey (email: j.blanden@surrey.ac.uk); Bratberg: Department of Economics, University of Bergen (email: Espen.Bratberg@uib.no); Gielen: Erasmus School of Economics (email: gielen@ese.eur.nl); van Kippersluis: Erasmus School of Economics (email: hvankippersluis@ese.eur.nl); Lee: Department of Health, Society, and Behavior, University of California, Irvine (email: keunbokl@hs.uci.edu); Machin: Department of Economics, London School of Economics (email: S.J.Machin@lse.ac.uk); Munk: The Physician Think Tank ATLAS, Copenhagen, and SEC, Uppsala University (email: mdmunk@gmail.com); Nybom: Institute for Evaluation of Labor Market and Education Policy (email: martin.nybom@ifau.uu.se); Ostrovsky: Statistics Canada (email: yuri.ostrovsky@statcan.gc.ca); Rahman: Frontier Economics (email: SumaiyaR@protonmail.com); Sirniö: Finnish Institute for Health and Welfare (email: outi.sirnio@thl.fi). Ilyana Kuziemko was coeditor for this article. We are grateful to Raj Chetty, Yonatan Berman, and three anonymous referees for feedback on previous drafts of this paper. An earlier version of this paper was presented at the ISA RC28 2019 Summer Meeting in Princeton, New Jersey. The authors gratefully acknowledge financial support from the following sources: the Harvard Multidisciplinary Program in Inequality and Social Policy (Manduca); the NORFACE Joint Research Programme on Dynamics of Inequality across the Life-Course, which is cofunded by the European Commission through Horizon 2020 under grant agreement no. 724363 (Adermon, Blanden, van Kippersluis, and Machin); NWO Vidi grant no. 452-17-007 (Gielen); the Strategic Research Council of the Academy of Finland, decision number 293103 (Sirniö); European Research Council Consolidator grant ERC-2013-CoG-617965 (Sirniö); and the Academy of Finland Flagship Programme, decision number 345547 (Sirniö).

† Go to https://doi.org/10.1257/app.20210137 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

increasing interest among scholars and policymakers in measuring absolute income mobility: the extent to which children grow up to earn more, in inflation adjusted terms, than their parents. Since 2016, studies have attempted to estimate rates of absolute income mobility in numerous countries around the world (e.g., Berman 2022; Bönke, Harnack, and Luthen 2019; Chetty et al. 2017; Chuard and Grassi 2020; Liss, Korpi, and Wennberg 2023).

Directly measuring absolute income mobility requires linked parent–child panel data extending across multiple decades. Because of these stringent data requirements, researchers have developed innovative methods for estimating population-level mobility rates based on more limited data (Berman 2022; Chetty et al. 2017). These approximation methods rely on a number of plausible but as-yet-untested assumptions about trends in relative income mobility and the shapes of the marginal income distributions among parents and children. In addition, cross-national absolute mobility comparisons have been challenging because estimated mobility rates may be sensitive to specification choices that vary across single-country studies due to data availability or local conditions. There has been little multinational analysis of the sensitivity of upward mobility patterns to substantively meaningful decisions about, for instance, the age at which to measure income or how to adjust for family composition.

In this paper we use linked administrative data from five countries in Europe and North America, as well as detailed survey data from two more, to address these two methodological challenges facing the rapidly growing literature on absolute income mobility. First, using linked administrative data from Canada, Finland, the Netherlands, Norway, and Sweden, we probe the assumptions behind the most common approximation method, in which cross-sectional data on the marginal income distributions of parents and children are combined with a rank–rank transition matrix, or copula, to estimate the average upward mobility rate for a given birth cohort.

For the five countries and the late twentieth-century birth cohorts covered by our administrative data sample, we show that this "copula and marginals" approximation method closely tracks the results produced from linked parent–child data. When using high-resolution, 100-cell income distributions (and a 100×100 cell copula), the point estimates never deviate by more than 1.6 percentage points from the true values, and definitive bounds on the measure that impose no further assumptions on the shapes of the marginal income distributions and copula never deviate by more than 5.6 percentage points. Even a much coarser, ten-cell distribution—of the type that could be constructed from longitudinal panel surveys such as the Panel Study of Income Dynamics (PSID) in the United States or the Socio-Economic Panel (SOEP) in Germany—produces point estimates within 3.0 percentage points of the linked-record values, though without further distributional assumptions the bounds on the estimate extend roughly 15 percentage points in either direction.

Our second central contribution is a multicountry evaluation of the sensitivity of upward mobility trends to substantively meaningful specification choices. We include seven countries in this portion of the analysis, adding the United Kingdom and the United States to the five countries mentioned above. In our baseline specification, upward mobility rates range from almost 80 percent for the Netherlands cohorts born in the 1970s to less than 50 percent for the 1984 US cohort. We evaluate the sensitivity of these results to several specification choices. First, we explore sensitivity to the age(s) at which parent and child incomes are measured. In doing so, we address concerns about the possibility of life-cycle bias—for instance, if recent cohorts of children reach full earnings potential later than their parents did—and concerns about the robustness of absolute mobility rates to short-term fluctuations in the business cycle or other transitory income shocks. For our full set of countries and cohorts, upward mobility rates are stable after age 35, and for Finland, the Netherlands, the United States, and the United Kingdom, they are stable after age 30. Mobility rates when income is averaged over several years are generally a few percentage points higher than those measured in one year, and show somewhat less fluctuation from year to year.

We also consider sensitivity to specification choices related to income concept, method of accounting for family structure, and price index. These decisions alter the substantive interpretation of results and do not necessarily have a "correct" answer. Absolute mobility rates measured with before-tax and after-tax income, for example, are both meaningful concepts, and either could be a legitimate object of study. For many of these specification choices, the size and even the sign of the effect differs from country to country, in ways that reveal aspects of their political, economic, and social environments. This heterogeneity highlights the need for care when comparing absolute income mobility rates across countries and for conceptual precision when choosing the income mobility concept to focus on, even in a single-country analysis.

In the following section, we review previous research on intergenerational income mobility, focusing in particular on recent scholarship measuring absolute mobility. We then provide an overview of our methods (more detailed descriptions of the methods used in each country are provided in online Appendixes 1 and 2). Next, we present our baseline estimates of absolute income mobility for our sample, and then the results of our validation of absolute mobility approximation methods and our multicountry sensitivity analysis. We conclude with a brief discussion of decomposition methods to pinpoint possible sources of variation in absolute mobility rates between countries and cohorts.

I. Prior Research

Scholars of intergenerational mobility distinguish between absolute mobility, which compares the raw outcomes of children and parents—in this case, their inflation-adjusted income—and relative mobility, which compares their rank or relative position in their respective distributions, adjusting for population-wide changes such as economic growth.

Historically, research on intergenerational income mobility focused on relative mobility: the association between the adult incomes of parents and children, often operationalized using the intergenerational elasticity of income or the correlation between the income rank of children and that of parents (Jäntti and Jenkins 2015; Torche 2015). A large literature has compared relative mobility rates across countries, generally finding that it is high in the Nordic countries and Canada, midrange in countries like Germany and Japan, and low in countries like Italy, the United

Kingdom, and the United States (Bratberg et al. 2017; Corak 2016; Harding and Munk 2020; Smeeding, Erikson, and Jäntti 2011).

Recently, a number of researchers have turned to absolute income mobility, motivated by the high salience of absolute comparisons among laypeople (Amiel and Cowell 1999; Ravallion 2018) and its straightforward normative interpretation: while one person's upward mobility in relative terms necessarily comes at the expense of someone else's downward mobility, upward mobility in absolute terms does not. Individual studies have estimated absolute mobility rates for recent cohorts of roughly 50 percent in the United States (Chetty et al. 2017), 53 percent in Canada (Ostrovsky 2017), between 59 percent and 68 percent in Australia (Deutscher and Mazumder 2022; Kennedy and Siminski 2023), 39 percent in Switzerland (Chuard and Grassi 2020), 70 percent in Germany (Bönke et al. 2019; Stockhausen 2018), and 77 percent in Sweden (Liss et al. 2023).[1]

A major spur to research on absolute income mobility has been the development of approximation methods that allow overall rates of absolute mobility to be estimated without linked parent–child panel data, by combining data on the marginal income distributions of parents and children with a copula, or transition matrix, which can be derived from another data source. Introduced to the absolute mobility context by Chetty et al. (2017), this method was expanded by Berman (2022), who adapted it for use without a copula and showed that absolute mobility rates are determined primarily by the marginal income distributions of parents and children, echoing findings for social class (Bukodi, Paskov, and Nolan 2019; Erikson and Goldthorpe 1992; Torche 2015). Applying his expanded method to ten high-income countries, Berman (2022) found evidence of widespread declines in upward mobility over the second half of the twentieth century.

While these approximation methods are very promising, they rely on several assumptions, described in detail below, that have not yet been directly tested by comparison to intergenerationally linked administrative data. Additionally, while the proliferation of absolute income mobility studies across countries raises the prospect of fruitful cross-national comparisons, research in individual countries has documented that upward mobility estimates can be quite sensitive to substantively meaningful specification decisions—for instance, about how to define income and how to adjust for family size (Chetty et al. 2017; Eshaghnia et al. 2021). There have not yet been multicountry studies using consistent definitions that evaluate the sensitivity of upward mobility rates to such specification choices.

There are reasons to expect that the countries in our sample might differ not only in their upward mobility rates but also in the way that absolute mobility rates are affected by particular specifications. Although all have market economies, they differ in their recent demographic histories and their economic institutions—for instance, in the size and robustness of their welfare states (Esping-Andersen 1990).

[1] Note that the usage of the terms "absolute mobility" or "absolute income mobility" in this paper and the literature described here is distinct from their usage by Chetty, Hendren, Kline, and Saez (2014), who use the term "absolute mobility" to refer to the expected adult income rank of a child born to parents at the twenty-fifth income percentile. That measure would capture a particular aspect of relative mobility in our terminology, since it is based on rank position rather than dollar earnings.

The remainder of this paper attempts to address these two outstanding methodological gaps in the absolute mobility literature. After providing an overview of the data and methods we use and presenting our baseline estimates of upward absolute income mobility by cohort for our sample countries, we conduct a detailed validation exercise of the "copula and marginals" approximation method for the countries and birth cohorts we observe. We then explore the sensitivity of upward mobility estimates to specification choices across all seven of our sample countries. We conclude with a discussion of methods for decomposing the difference in mobility rates between two countries or cohorts into contributing factors.

II. Overview of Methods and Data

Absolute income mobility measures the extent to which children grow up to earn higher inflation-adjusted incomes at a given age than their parents did. Following Chetty et al. (2017), for a particular cohort of children at age a, we calculate the upward absolute income mobility rate A_a as

$$A_a = \frac{1}{N}\sum_{i=1}^{N} U(p_{ia}, c_{ia}),$$

where

$$U(p_{ia}, c_{ia}) = \begin{cases} 1, & \text{if } c_{ia} \geq p_{ia} \\ 0, & \text{if } c_{ia} < p_{ia} \end{cases},$$

N is the number of children in the cohort, p_{ia} is the inflation-adjusted income of the parents of child i when they were age a, and c_{ia} is the inflation-adjusted income of child i at age a.

We use two primary methods to measure absolute income mobility. The more straightforward but data-intensive method is to calculate absolute mobility directly from intergenerationally linked records. For our sample countries where register data that link children to parents and track incomes over time are available—Canada, Finland, the Netherlands, Norway, and Sweden—we measure the household incomes of children and their parents when each is age 30.[2] We then adjust for inflation using each country's consumer price index (CPI) and compute the fraction of children whose incomes match or exceed their parents'. This is the most straightforward way to measure absolute mobility, requiring minimal assumptions or statistical approximations, but it is only possible where large-scale survey or administrative panel data exist.

The second method we use is the "copula and marginals" approach introduced by Chetty et al. (2017). This involves constructing a copula, or parent–child income rank transition matrix, for the subset of the data where linked income information is available for parents and children. We use both 10×10 decile cell and 100×100 percentile cell copulas in our validation exercises. Our baseline

[2] As described in online Appendix 2, in some cases where income is not observed at age 30, we use income at the observed age closest to 30 within a certain range, typically five years.

estimates for the United Kingdom use a 10×10 copula, while those for the United States use a 100×100 version. We then create marginal income distributions for parents and children in each birth cohort. We calculate the overall absolute mobility rate for a given cohort by first comparing the mean incomes in every pair of child and parent cells and determining whether the children in that child cell had higher incomes than the parents in that parent cell. Next, we take the average of upward mobility across all parent and child cell pairs, weighting by the probability from the copula that children with parents in that parent cell grew up to have incomes placing them in that child cell. This approach does not determine whether any individual child outearned his or her parents, but it does provide an accurate estimation of the upward mobility rate in total, as we show below.

Register data for Canada, Finland, the Netherlands, Norway, and Sweden were accessed via the national statistical agencies of those countries (Statistics Canada 2019; Statistics Finland 2020; Statistics Netherlands 2020; Statistics Norway 2018; Statistics Sweden 1800–2016, 1932–2016). For the United Kingdom, we combine data from the Family Expenditure Survey (Central Statistical Office 1997; Department of Employment 1990; Office for National Statistics 2002), the British Cohort Study (Bynner, Butler, and University of London, Institute of Education, Centre for Longitudinal Studies 2022; University of London 2013), and the Family Resources Survey (Department for Work and Pensions, Office for National Statistics, NatCen Social Research 2018). For the United States, we construct parent and child marginal income distributions from the decennial census, American Community Survey, and Current Population Survey, accessed via IPUMS (Flood et al. 2022; Ruggles et al. 2022). The US copula is sourced from Chetty et al. (2017) and derived from anonymized income tax records (Chetty, Hendren, Kline, and Saez 2014).

Because the type, time period, format, and quality of data differ across the countries in our sample, the data and methods that we use vary somewhat from country to country. A high-level comparison of the specifications used across all seven countries is provided in online Appendix 1, and detailed descriptions of the exact data, methods, and specifications used in each country are provided in online Appendix 2.

Our baseline specification, which allows the closest possible match in the largest number of countries and cohorts, defines income as the pretax sum of income from (i) the labor market, (ii) self-employment earnings, and (iii) social insurance programs such as unemployment insurance and social security. We sum the individual incomes of the target child and his or her spouse (if present) at age 30 and compare them to the sum of the individual incomes of the child's parents when they were 30, adjusted using the country's official CPI. If the parents are not the same age as one another, we measure their income when the father was 30 if a father is present and measure when the mother was 30 otherwise.[3]

[3] Data constraints require us to deviate from these exact definitions in certain countries. As described in online Appendixes 1 and 2, in Finland, we include capital income as well as labor and transfer income, while in Canada, the Netherlands, and the United Kingdom, we include income from social transfer programs beyond unemployment and social security. In Norway, we measure parent age using the father only, while in the United Kingdom, we measure it using the parent who results in the higher total income.

We limit our sample to children born in the country everywhere except Canada, where children are included if they lived in Canada between ages 16 and 19; Finland, where children are included if they lived in the country at age 30 and can be linked to their parents in the register data; and the United Kingdom, where children are included if they were present in the country at age 30. Because all of the children in a given cohort were by definition born in the same year, their incomes are measured in the same year as one another, 30 years later. But since parents vary in the age at which they have children, parent incomes for a given cohort will include data from more than one calendar year. One point to note is that because we compare children to their parents, our analysis compares all members of each child cohort to only those members of the parent cohort who had children. If there is selection into parenthood, our results for absolute mobility may depart from trends in average income across generations.

Except in the United Kingdom, we include both parents and children with zero income in our sample as long as they are observed in the data (if they have died or left the country, they may not appear). Because we define children as upwardly mobile if their incomes are at least as great as their parents', children of parents with zero income will be upwardly mobile as long as their incomes are nonnegative. Unlike many measures of relative mobility, which use correlations that can be substantially affected by the inclusion of zeros, the impact of zero incomes on our measure of absolute mobility is limited to only those individuals who had zero income themselves.

In alternative specifications discussed in detail in Section V, we vary (i) the age of income measurement, (ii) the number of years of income data used, (iii) the income sources included, (iv) the method of adjusting for different family structures, and (v) the price index used to adjust for inflation.

III. Baseline Estimates of Absolute Income Mobility

Before beginning our detailed methodological discussion, we briefly present our baseline estimates of absolute income mobility for our sample countries and cohorts. Figure 1 displays trends in absolute income mobility by birth cohort for the countries included in this study. Because of data limitations, not all countries have estimates for all cohorts. The US series goes back to the 1940 birth cohort, while most European countries begin in the 1960s. Data for Canada only exist for the 1976–1985 birth cohorts, and for the Netherlands we focus on the 1973–1984 cohorts. For ease of comparison, we show results starting in 1960, the first year for which non-US data are available.

As Figure 1 shows, there was substantial variation across these countries in both the level of absolute income mobility and the trends over time. At the top end, recent cohorts of Norwegians have experienced upward mobility rates of roughly 75 percent, while cohorts born before 1980 in the Netherlands saw upward mobility rates approaching 80 percent. The United States, Canada, and Finland show lower rates of upward mobility, under 60 percent, and declining trends over time. Summary statistics for the earliest and latest cohorts in each country are shown in online Appendix 3. Some important points to note are that the average parent age at

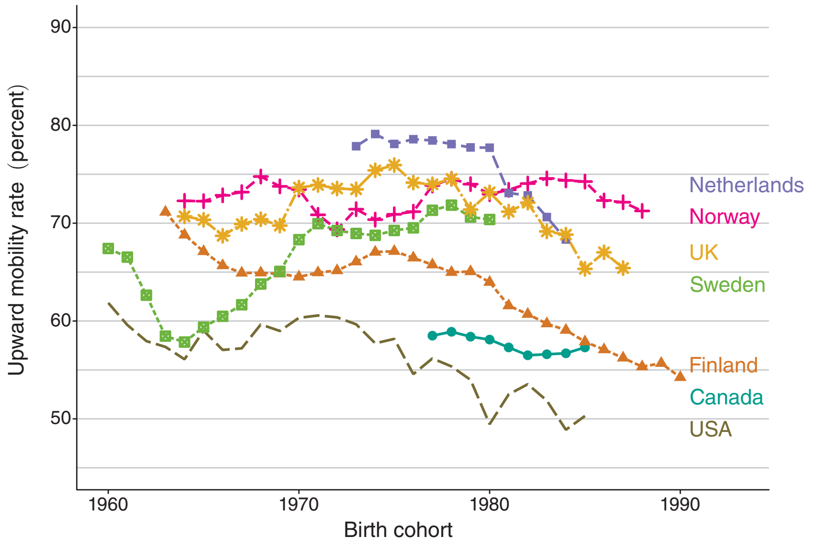

Figure 1. Baseline Estimates of Upward Absolute Income Mobility by Country and Birth Cohort

Notes: The upward mobility rate is calculated as the percentage of children in each birth cohort whose family income at age 30, adjusted for inflation, was at least as high as their parents' family income at age 30. Incomes are measured using a combination of register and survey data in each country, as described in online Appendix 2. Parent ages are defined using the methods described in the main text and online Appendixes 1 and 2.

childbirth increased substantially during our period in most sample countries, and that rates of marriage and/or cohabitation were generally higher among parents than among children. This largely reflects changing patterns of marriage and household formation but in certain cases (most notably Norway and Sweden, as discussed in online Appendixes 2.4 and 2.5 respectively) may also reflect data artifacts. Also of note is that the fraction of children in school at age 30 has risen over time in many of our sample countries. To the extent that those students also have low incomes, this may be a reason that upward mobility measured at age 30 is sometimes lower than mobility measured at older ages, as shown in Figure 4.

Relative income mobility, as measured using rank–rank slopes (Chetty, Hendren, Kline, Saez, and Turner 2014), ranged from 0.07 in Norway for cohorts born in the mid-1970s to 0.34 for the US 1980–1982 cohorts. Within the countries where we are able to observe relative income mobility over time, there was generally fluctuation in relative mobility rates but not secular increases or decreases—with the exception of the Netherlands, where the rank–rank slope fell from around 0.23 for the 1973 birth cohort to 0.16 for the 1984 cohort. Online Appendix 4 makes available all 76 100×100 and 96 10×10 empirical copulas that we construct and displays trends in the rank–rank slopes over time.

We provide these copulas primarily for others who may wish to use them in bounding analyses of absolute mobility as described in Section IVA. Caution should be used when comparing these data to previous studies of relative mobility because they are constructed for a single year of income at age 30, creating the potential for life-cycle and errors-in-variables bias, both of which would bias estimates downward, i.e., in the direction of higher relative mobility (Mazumder 2005; Nybom and Stuhler 2017; Solon 1992). That said, the generally stable trends we observe in the Nordic countries are consistent with previous research (Engzell and Mood 2021; Pekkarinen, Salvanes, and Sarvimäki 2017; Sirniö, Kauppinen, and Martikainen 2017).

In online Appendix 5 we present a detailed comparison of our results with those of Berman (2022). The upward mobility rates we calculate for recent cohorts roughly align with Berman's (2022) estimates in most countries but differ in the United Kingdom by as much as 10 to 15 percentage points. Additionally, we find stability in absolute mobility rates over time in certain countries where Berman (2022) reports downward trends. Where differences exist, they appear to be due to differences in the data used rather than differences in methods—specifically, the fact that we measure incomes for 30-year-old parents and children directly, while Berman (2022) uses income data for the entire population. Because younger workers tend to be more affected by changes in the macroeconomic climate than older workers (Hoynes, Miller, and Schaller 2012), and because incomes of parents may differ systematically from incomes of nonparents of similar age, trends in the full income distribution are in some cases not an accurate proxy for trends in the earnings of young adults. As discussed in Section IVC, this substitution of not-quite-identical distributions is one of the most likely—and most difficult to identify—sources of error when using the "copula and marginals" approximation.

IV. Validation of the "Copula and Marginals" Approach

Recent scholarship on trends in absolute mobility over time has been facilitated by the development of approximation methods to estimate upward mobility rates in the absence of linked parent–child data. Chetty et al. (2017) proposed that overall rates of absolute mobility can be accurately calculated without linked panel data, by combining data on the marginal income distributions of children and parents with the copula, or parent–child rank transition matrix. This approach draws on Sklar's theorem (Sklar 1959), which showed that any multivariate distribution can be expressed in terms of marginal distributions and a copula. Because of its much lower data requirements and its ability to incorporate data from multiple sources, the "copula and marginals" approach is becoming widely used in studies of absolute income mobility (e.g., Berman 2022; Bönke et al. 2019).

The logic behind this approach is compelling, but it has never been validated through a direct comparison of absolute mobility rates calculated using the "copula and marginals" approach and those using the true, linked-records approach on the same data. We conduct such a comparison here. For four of the countries in our sample—Canada, Finland, Norway, and Sweden—we can calculate the upward mobility rate directly from linked data and also produce copulas and marginal distributions. By comparing estimates constructed using the "copula and marginals" approach to those constructed from linked data, we are able to determine whether the former is a reasonable approximation of the latter.

While Sklar's theorem is a mathematical identity, the way it is operationalized to estimate absolute mobility rests on a number of assumptions that, if they fail to hold, could potentially introduce discrepancies between the "copula and marginals" and linked-records approaches. In practice, we identify three primary possible sources of discrepancy. First, if the copula is constructed from only a subset of cohorts, it might not reflect the true transition probabilities between parent and child income ranks for all considered cohorts. Second, the process of discretizing the (essentially) continuous marginal income distributions of parents and children into percentile or decile cells, and assigning the same mobility outcome to all children in a cell, may add imprecision to the estimates. Finally, if the marginal income distributions are approximated using data from a population other than that of cohort parents and children, such as the entire population of a country, that could bias the overall estimate in ways that may be difficult to detect.

This section considers each of these possible sources of discrepancy in turn, for both a high-resolution 100×100 copula (the type used by Chetty et al. 2017, which requires population-level administrative data for at least one cohort to create) and a lower-resolution 10×10 copula (which is possible to construct from longitudinal surveys such as the US PSID or German SOEP). We also consider even more simplified methods of estimating upward mobility by comparing just the median incomes of parents and children or by calculating the fraction of children outearning the median parent (Katz and Krueger 2017).

For the countries and cohorts in our sample, we show that the assumption of copula stability is not a major source of potential error for the "copula and marginals" approach. More concerning is the process of discretization, which does not substantially bias the point estimates of mobility but does extend the assumption-free bounds on the estimates quite substantially in either direction, particularly for the 10×10 copula. Potentially most problematic is the third source of discrepancy, the use of alternate income distributions to stand in for the marginal distributions of children and parents. Depending on patterns of income in each country, this substitution can potentially induce considerable bias into the results. For this reason, we recommend that researchers do as much as possible to validate that the data they use to construct their marginal income distributions are, in fact, a good approximation of income data for the exact population being studied.

A. *Assuming Copula Stability Is Unlikely to Be a Major Source of Error*

The first source of potential error we consider is the assumption that one specific copula—often generated from a recent cohort where linked parent–child data are available—can be used to construct upward mobility estimates for earlier cohorts. This assumption will be unproblematic if either (i) the true copula (i.e., the rate of relative income mobility) is fairly stable over time or (ii) which copula is used has minimal impact on the calculated upward mobility rate.

Figure 2 plots the true, linked-records rate of upward absolute mobility (solid line) as well as two versions of the estimated rate using the "copula and marginals" approach with high-resolution 100×100 copulas. The dotted line shows the "copula and marginals" estimate using a copula constructed based on data from the exact cohort being estimated, while the dashed line uses the most recent copula for all cohorts in a given country. Both estimates hew closely to the true rate: across all four

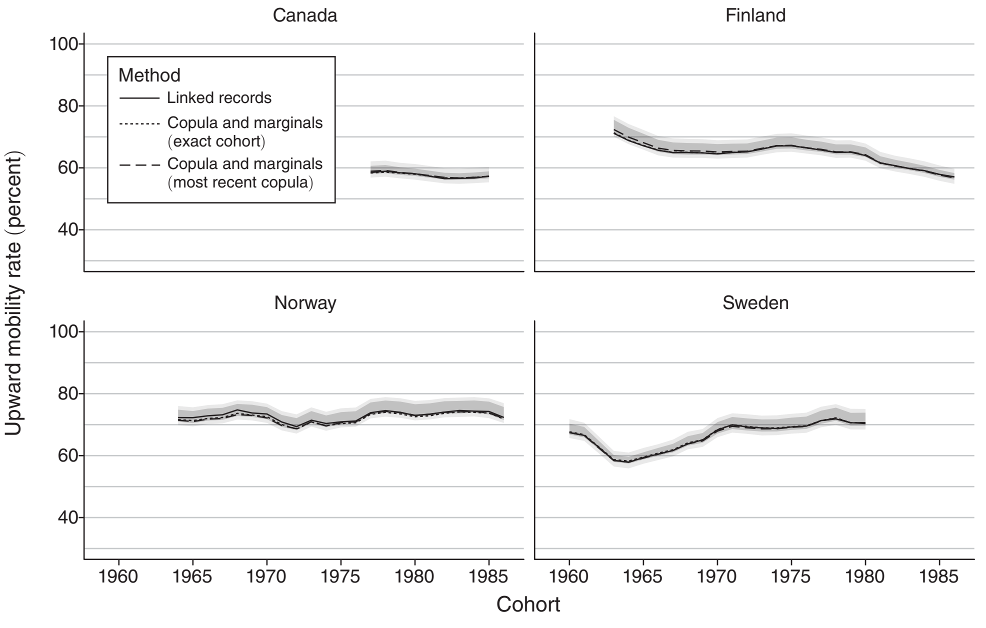

FIGURE 2. VALIDATION OF THE "COPULA AND MARGINALS" APPROACH FOR ESTIMATING ABSOLUTE INCOME MOBILITY, 100×100 COPULA

Notes: This figure compares estimates of absolute income mobility using linked records with those using the "copula and marginals" approach introduced by Chetty et al. (2017). "Copula and marginals" estimates are constructed by computing child and parent marginal income distributions at age 30 for each birth cohort in each country and combining them with the 100×100 percentile cell parent–child income rank transition matrix constructed based on linked parent–child records in each country. Dark gray shading shows bounds across all empirical copulas we observe. Light gray shading shows bounds with no assumptions about the within-cell shapes of the marginal distributions.

countries for which both approaches are implemented, the exact copula estimate never varies by more than 1.13 percentage points from the true value, and the most recent copula estimate never varies by more than 1.53 percentage points. For comparison, the true rate of upward mobility changed by 2.4 percentage points in Canada during our sample period and by 16.9 percentage points in Finland.

Going one step further, the area shaded in dark gray shows the range of estimates generated when we apply all 76 100×100 empirical copulas that we observe across countries and birth cohorts in our sample to the marginal distributions for each country-cohort. Even though rates of relative income mobility varied substantially within our sample—the parent–child rank–rank slopes we observe range from 0.07 for the 1976 cohort in Norway to 0.34 for the 1980–1982 cohorts in the United States—the bounds constructed by considering all empirical copulas never extend more than 4.36 percentage points from the true value. The highest estimates generally result from the 1980–1982 US copula, which had the lowest relative mobility rate, consistent with the predicted inverse relationship between relative and absolute mobility (Berman 2022). Because of this, while we have made all 76 empirical copulas available to future researchers in online Appendix 4 for the construction of copula bounds in other studies, those working in countries with lower rates of relative

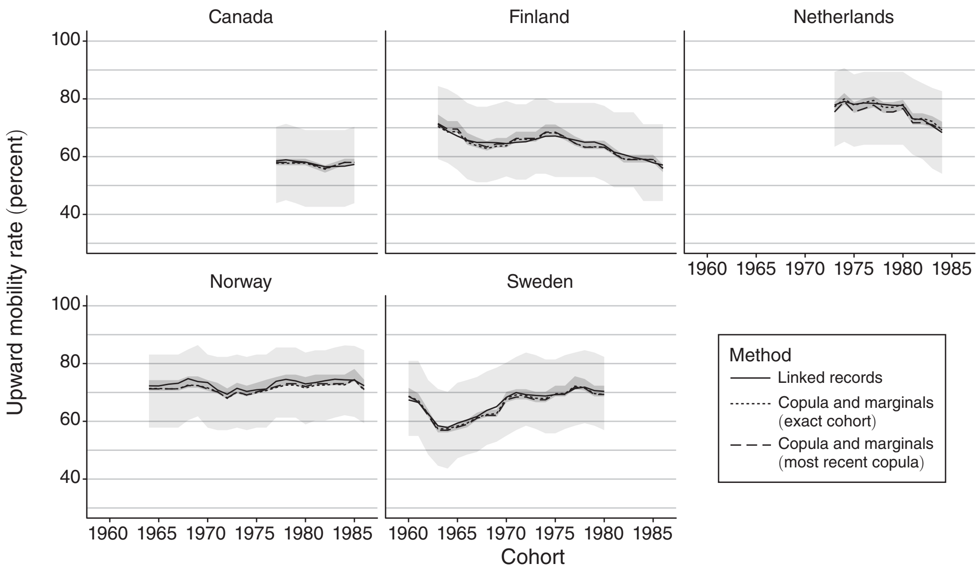

Figure 3. Validation of the "Copula and Marginals" Approach for Estimating Absolute Income Mobility, 10×10 Copula

Notes: This figure compares estimates of absolute income mobility using linked records with those using the "copula and marginals" approach introduced by Chetty et al. (2017). "Copula and marginals" estimates are constructed by computing child and parent marginal income distributions at age 30 for each birth cohort in each country and combining them with the 10×10 decile cell parent–child income rank transition matrix constructed based on linked parent–child records in each country. Dark gray shading shows bounds across all empirical copulas we observe. Light gray shading shows bounds with no assumptions about the within-cell shapes of the marginal distributions.

income mobility than the United States should attempt to generate additional copulas from their own data to provide more accurate upper bounds.

Figure 3 displays the same analysis using the coarser, 10×10 copula. As with the 100×100 version, the point estimates using both the exact and the most recent copulas closely track the linked-records values—in this case, never deviating by more than 2.40 percentage points using the exact copula or 3.01 percentage points using the most recent copula. The bounds constructed from all empirical copulas never extend more than 5.19 percentage points from the true value. We are able to display the Netherlands in Figure 3 but not Figure 2 because our linked-records sample there is large enough to create a 10×10 copula but not a 100×100 copula.

Overall, this exercise gives us confidence that the assumption of copula stability presents relatively little risk to accuracy of the "copula and marginals" approach—at least for these countries and birth cohorts. While the country-level copulas we observe are fairly stable over time, as shown in online Appendix 4, the bounds are narrow even for the estimates that are generated from applying quite different copulas (for example, applying the US copula to Norway or Finland). This suggests that our result may be driven primarily by the relatively limited impact that variation among empirically observed copulas has on rates of absolute mobility, as previous research has argued (Berman 2022).

B. *Discretizing the Marginal Distribution Expands Bounds but Not Point Estimates*

A second assumption introduced by the "copula and marginals" approach has to do with the way in which an (almost) continuous income distribution among the population of children and parents is approximated by a discrete 100-cell or 10-cell marginal income distribution. As typically applied, the method compares the mean income in each child cell to that in each parent cell, assigning all children in that cell to be upwardly mobile if their mean exceeds the mean of the parent cell. This method is correct when there is no overlap between the income ranges for children and parents in a particular cell—if all children in a cell earn more than all parents, it must be the case that upward mobility for that cell is 100 percent. However, in cases where the two cell distributions overlap, it is likely that some individual children in that cell will be upwardly mobile, while others will be downwardly mobile. Assigning all of them a mobility outcome based on the cell mean obscures this detail. We believe that this "discretization error" may be responsible for instances—particularly with the 10×10 copula—where the linked-records mobility rate falls outside the bounds created by applying all empirical copulas.

The light-gray shading in Figures 2 and 3 shows the range of upward mobility rates that are consistent with the parent and child marginal distributions after relaxing the assumption that the comparison of mean cell values correctly identifies upward mobility in cells whose distributions overlap. The lower bound is constructed by assuming that all children in cells where the parent and child distributions overlap grow up to earn less than their parents, and applying the empirical copula that gives the lowest upward mobility rate under those conditions. The upper bound is constructed by assuming that all children in overlapping cells earn more than their parents, and applying the empirical copula that gives the highest rate under those conditions. These bounds thus make no assumptions about the shapes of the parent and child marginal distributions beyond the observed bounds of each cell, and assume only that the true relative mobility rate is within the range generated by the empirical copulas we observe.

When using the 100×100 cell copula, shown in light gray in Figure 2, the bounds added by relaxing the assumptions related to discretization broaden the possible upward mobility somewhat, but still within a relatively narrow range—roughly 5 percentage points above or below the true value. In contrast, results for the 10×10 copula, shown in Figure 3, show bounds that extend roughly 15 percentage points in either direction. With these wider bounds, it is much more difficult to have confidence about the relative positions of different countries or the extent of changes over time.

Unless accompanied by assumptions about the shape of the copula and marginal distributions, the process of discretization introduces a potentially large amount of error into the "copula and marginals" estimates—as much as 5 percentage points when using a 100×100 copula, and 15 percentage points when using the lower-resolution 10×10 copula. In our sample countries and cohorts, the use of mean cell values did not result in error this large, but that possibility cannot necessarily be ruled out elsewhere. For this reason, even though the point estimates from 10×10 copulas do appear to generally track the true upward mobility rates in our sample, we recommend using higher-resolution copulas and marginal distributions when possible to narrow the assumption-free bounds. Our own estimates of upward mobility in the United Kingdom, which rely on a 10×10 copula due to the small sample size available, should be interpreted with some caution for this reason. The good news is that as long as the cutoff values for each cell in the marginal distributions are known, it is possible to precisely calculate the bounds on the estimated absolute mobility rate that arise from this source of potential error.

C. *Representativeness Error*

The final source of potential error in the "copula and marginals" analysis that we identify can arise when the data used to create the marginal income distributions of parents and children in a given cohort are not, in fact, representative of that cohort. Of the three sources of potential error, we believe this is the most concerning, because it is challenging to predict, potentially large in magnitude, and cannot be identified or bounded without access to the true data.

Probably the most common circumstance under which this error will arise is when income data for the entire population are used to stand in for that of a specific cohort of children and parents. Data on the income distribution of the entire population may be available in countries where data on income by age are not. In many cases, income distributions and trends for young adults will closely track those for the population overall. But in some cases, they do not.

We believe that this lack of representativeness is the primary reason for discrepancies between our baseline estimates and those of Berman (2022), who used a variation on the "copula and marginals" method to estimate absolute mobility rates in several of our sample countries. As we describe in online Appendix 5, for most sample countries, our estimates track Berman's (2022) quite closely, but in certain cases—the United Kindom in particular—they diverge markedly. Detailed examinations of these cases in online Appendix 5 of this paper and online Appendix C of Berman (2022) suggest that this divergence is due to differences in the income trends of 30-year-olds compared to the entire population in those countries. Unfortunately, it may not always be possible to determine the extent to which population-wide income trends are accurate indicators of trends for particular age groups. In online Appendix 6, we compare the median income among child cohorts in our sample countries with the median household income (excluding capital income and social transfers) among 30-year-olds for similar cohorts measured in the Luxembourg Income Study (2023). In Finland, the Netherlands, Norway, and Sweden, the trends match closely, while in Canada, the United Kingdom, and the United States, the median incomes among our sample children are lower than those in the LIS.

Even when income data by age are available, they might not be representative of children and parents. One conceptual quirk of absolute income mobility is that while the children in a given cohort closely (ideally exactly) match the entire population of people born in a particular country in a particular year, the parents of that cohort do not match the population born in any one year or even any weighted combination of years. Not everyone has children, and if there is selection into parenthood by income, that may mean that parent incomes diverge from the incomes of people born in the same period who are not parents. While this would not necessarily create inaccuracy in the absolute mobility measure, since the goal of the measure is to compare children to their parents, it could create a divergence between absolute mobility and more general trends in income levels for the entire population. If the amount of selection by income varies over time, this could cause fluctuations in the degree to which absolute mobility rates track average income growth; for example, if parent incomes have grown relative to nonparent incomes, this would likely lead to lower rates of absolute mobility than trends in average incomes would predict. Future research should explore the extent to which selection into parenthood shapes absolute mobility rates.

D. *Median Income Approximation*

Katz and Krueger (2017) proposed even more ambitious methods of approximating trends in absolute income mobility. Noting that the upward mobility rates calculated by Chetty et al. (2017) were correlated at 0.995 with the inflation-adjusted difference in median incomes between parents and children, they suggested that it might be possible to reliably approximate upward mobility rates using the difference in median incomes alone. They also suggested approximating absolute mobility with the fraction of children whose incomes exceed those of the median parent. These methods would further relax the data requirements for calculating absolute mobility rates.

We explore these possibilities in online Appendix 7. Ultimately, we believe that the methods proposed by Katz and Krueger (2017) do an excellent job of capturing whether or not economic conditions are improving over time, both from parents to children and from one cohort to the next. But we hesitate to recommend them as a means of constructing a precise estimate of absolute mobility that could be compared to estimates from the linked-records or "copula and marginals" approaches in other countries or periods. Considering first the parent–child difference in median incomes, while the strikingly high correlation that Katz and Krueger (2017) observed in the United States also appears in several other countries, it is much lower for others—most notably Norway, where the correlation is just 0.62. Additionally, the process of translating from parent–child median income differences measured in local currency units to absolute mobility rates measured in the proportion of children outearning their parents will require a different translation factor for each country, which the researcher must choose.

The correlations between absolute mobility and the fraction of children outearning the median parent are also quite high (above 0.96) for Finland, the Netherlands, Sweden, and the United States, but they are somewhat lower, around 0.9, for Canada and Norway. Unlike the median income difference, this measure does require data on the full income distribution of children, though not that of parents.

Because the correlations are not consistently high across countries, for countries and cohorts where median income data are available but full distributions are not, we recommend simply reporting changes in median incomes over time and perhaps referencing them against similar changes in countries where absolute mobility can be calculated.

V. Cross-National Analysis of Sensitivity to Specification Choices

The previous section investigated the validity of the increasingly common "copula and marginals" approximation of absolute income mobility. In that analysis, there was a clear correct answer—the upward mobility rate calculated from linked parent–child data—and the central question was how closely the "copula and marginals" approach came to the correct result and what conditions might prompt it to fail. We now turn to a second methodological question: how absolute mobility rates are affected by substantively meaningful specification choices, and how the effect might vary from country to country. Specifically, we consider decisions related to the age and number of years when parent and child incomes are measured, the income sources that are included, the way family structure is taken into account, and the price index used.

Unlike in the previous section, here there is not a clear correct answer: two different specifications will each produce valid estimates of absolute mobility, albeit ones with distinct substantive interpretations. As we show, specification choices do impact the calculated absolute mobility rates, sometimes substantially, and the size and even sign of the effects vary from country to country. Because of this, we recommend that researchers think carefully about which specification best captures the concept of absolute mobility they are interested in, and do their best to match specifications when making comparisons across countries.

We first explore the sensitivity of upward mobility estimates to the age of income measurement and the number of years of income that are observed. We then consider sensitivity to income concept, family structure, and price index.

A. *Sensitivity to Age of Income Measurement and Number of Years of Income Observed*

Studies of relative income mobility have long contended with two challenges related to the age of income measurement. First, age–income profiles might vary over time and among members of the same cohort—for instance, if some individuals stay in school longer than others (Grawe 2006). Second, earnings may fluctuate from year to year. Each of these phenomena means that relative mobility rates calculated for a particular cohort based on income data from a particular year may not necessarily be indicative of rates calculated from other income data, even for the same cohort (Nybom and Stuhler 2016).

In contrast to relative income mobility, which is typically operationalized as the slope of a regression line of child on parent earnings, absolute income mobility is not likely to be affected by attenuation bias due to earnings fluctuations from year to year. However, because it is measured using dollars rather than ranks, absolute income mobility rates may be especially sensitive to macroeconomic trends, including the business cycle. A widespread recession will only impact estimates of relative mobility to the extent that it impacts some people more than others and therefore reshuffles income ranks. But even a rank-preserving recession can have a very large impact on the overall rate of absolute mobility. The early 1990s recession in Sweden, for example, is likely a major reason for the abrupt drop in the Swedish

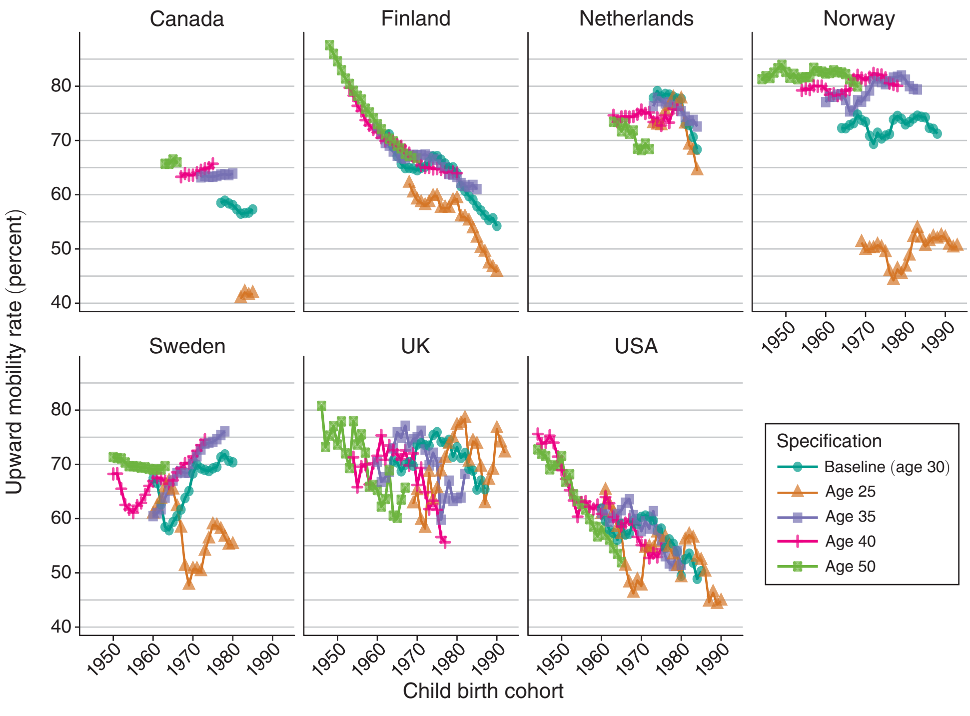

Figure 4. Sensitivity of Absolute Income Mobility Rate to Age of Income Measurement

Notes: This figure displays upward mobility rates calculated using a single year of income measured at ages 25, 30 (baseline), 35, 40, and 50 for each country in our sample. Parent ages are defined using the methods described in the main text and online Appendixes 1 and 2.

upward mobility rates shown in Figure 1 for cohorts born in the early 1960s (whose adult incomes were measured in the early 1990s), followed by an equally abrupt rise for cohorts born in the late 1960s.

Here we explore the sensitivity of upward mobility estimates to decisions about the age at which income is measured and the number of years of income that are incorporated into the estimates. As mentioned, there is no definitively correct age or age range at which to measure absolute mobility—while earning more than one's parents did at age 30 means something different than the same achievement at age 40, both can reasonably be considered absolute mobility. However, for many purposes it may be desirable to calculate mobility in permanent income, a person's mean income averaged over their entire life. This cannot be perfectly measured until both parents and children are deceased, but if mobility rates stabilize after a certain age or when averaging more than a certain number of years, that could suggest a workable approximation.

Age of Income Measurement.—We first consider the extent to which rates of absolute income mobility systematically vary with the age at which income is measured. For this analysis we compare upward mobility rates calculated using one year of income data measured at the following ages: 25, 30 (baseline), 35, 40, and 50. The results are shown in Figure 4. Two main patterns stand out. First, in several of the countries—Canada and Norway, and to a lesser extent Sweden and Finland—there does seem to be a systematic association between age of income measurement and absolute mobility, especially below age 35. In Norway, for example, upward mobility rates measured at age 25 are around 50 percent, while those at age 30 (our baseline) are around 70 percent, and at ages 35 and over, they stabilize around 80 percent. In Canada and Sweden, the gap between ages 25 and 35 averages around 15 percentage points for the cohorts where both ages are observed, while in Finland it averages around 7 percentage points. Interestingly, the Netherlands, United Kingdom, and United States do not appear to have a strong relationship between upward mobility rates and the age of income measurement, perhaps indicating that parents and children in these countries have more stable age–earnings relationships. Still, given these results, we advise researchers to avoid using income data from ages below 30 and ideally to use data for ages 35 and up.

A second pattern in Figure 4 is the impact of the business cycle and macroeconomic shocks on upward mobility rates. This is most clear in the graph for Sweden, where the trough in the age 30 mobility rate for the early 1960s cohorts is repeated in the age 25 results for the late 1960s cohorts, the age 35 results for the late 1950s cohorts, and the age 40 results for the early 1950s cohorts—all of which had their incomes measured during the early 1990s recession. A similar pattern can be seen in Norway with the slight dip between the 1970 and 1975 cohorts in the baseline analysis that is repeated five cohorts later in the age 25 analysis and five cohorts earlier in the age 35 analysis. To a lesser extent, the United Kingdom and United States show the impacts of the early 2000s and late 1990s booms, respectively.

Number of Years across Which Income Is Averaged.—The fact that the business cycle can be seen so clearly in the single-year upward mobility estimates in Figure 4 should give researchers pause. While mobility rates calculated from a single year of income are accurate statements for that particular year, they may not give an accurate portrayal of a person's overall mobility trajectory, especially if economic conditions in the year of measurement were atypical. Figure 5 assesses this possibility, as well as the sensitivity of absolute income mobility rates to transitory income fluctuations. It plots three series: absolute income mobility estimated with a single year of data at age 35, that estimated using the mean income from ages 33 to 37, and that using the mean of ages 30–40. We use age 35 rather than 30 as our central point to avoid incorporating incomes before 30 into the averages, in light of the findings from Figure 4.

For the countries with sufficient years of data, the multiyear averages tend to show upward mobility rates that are a few percentage points higher than the single-year values for the same cohorts, and that have fewer fluctuations from year to year. This is most visible in Norway, where the dip in upward mobility for the 1965–1970 cohorts that appears in the single year age 35 series is reduced in the 5-year age 33–37 series, and disappears almost entirely in the 11-year age 30–40 series. Interestingly, the dip and subsequent rebound that is visible for the Swedish 1958–1963 cohorts in the

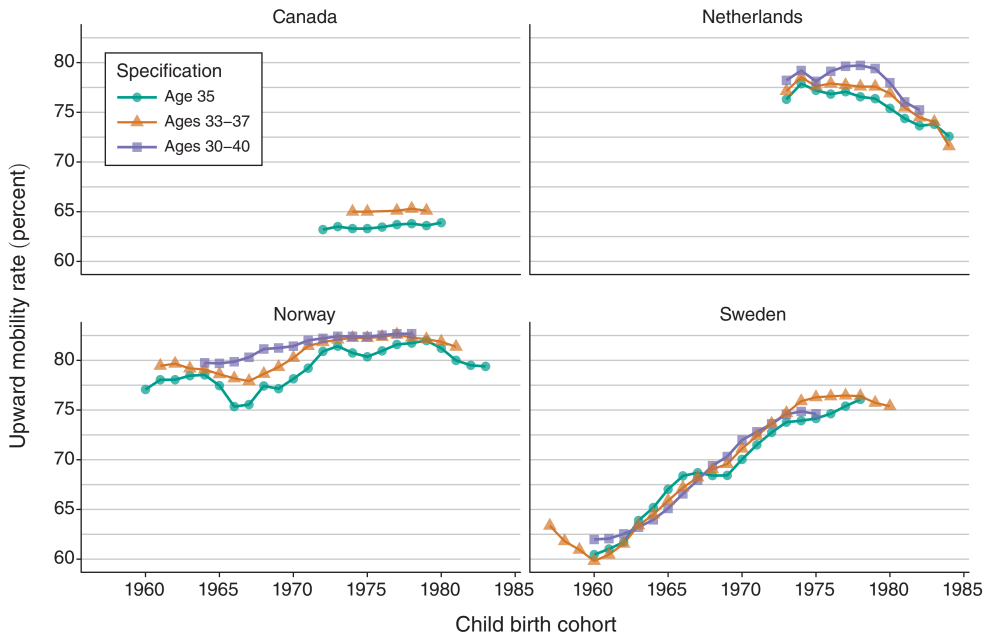

Figure 5. Sensitivity of Absolute Income Mobility Rate to Averaging Income across Multiple Years

Notes: This figure displays upward mobility rates calculated by averaging income across 1, 5, and 11 years for the countries in our sample where data permit. Parent ages are defined using the methods described in the main text and online Appendixes 1 and 2.

single-year data appears generally similar in the more extended series, perhaps indicating a longer-term recession.

Where possible, we encourage researchers to include multiple years of income data when calculating absolute mobility if their goal is to approximate permanent income. However, single-year estimates do appear to track multiyear averages fairly closely, so they may be suitable substitutes when multiple years are not available.

B. *Sensitivity to Income Concept, Family Structure, and Price Index*

We also explore the sensitivity of absolute mobility estimates to specification choices about the income concept, method of accounting for family structure, and price index. Even more than results at different ages, these specifications have distinct substantive interpretations from one another. The differences between the specifications considered here also reveal aspects of the distinct political, economic, and social conditions of the countries in our sample.

Income Concept: Disposable versus Pretax Income.—In our baseline analysis, we define income as the sum of labor earnings, self-employment income, unemployment insurance, and pension benefits, measured before taxes. This is the income definition that is most consistently available across the full range of countries in

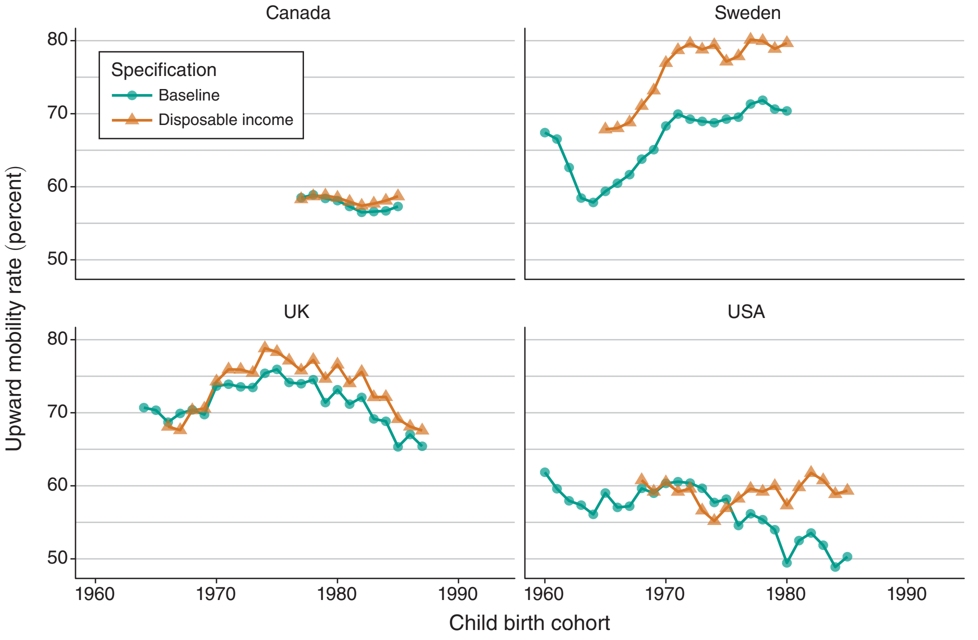

Figure 6. Disposable Income Analysis

Notes: This figure compares our baseline estimates with those using posttax disposable income in the countries where both measures are possible. In Canada and the United Kingdom, the baseline specification includes social transfers.

our sample. But it omits several potentially major sources of income, most notably capital income and income from social transfer or income support programs, such as the Earned Income Tax Credit or Temporary Assistance for Needy Families in the United States. It also does not account for taxes. Posttax disposable income is generally thought to offer a better measure of true living standards than pretax income (Canberra Group 2011).

In Figure 6 we plot absolute mobility rates using disposable income, defined as income from all sources (including capital and social transfer programs), after taxes. This series is only available for four of our seven sample countries, and is shown in the orange triangles in Figure 6. As the figure shows, in Canada and to some extent the United Kingdom, upward mobility in disposable income closely tracks that in our baseline, pretax specification, but in Sweden and for recent US cohorts, upward mobility in disposable income is substantially higher, by almost 10 percentage points. It is important to note that the baseline specifications in Canada and the United Kingdom include social transfer programs, which may be the reason for the smaller differences in those cases. These results suggest that the tax system, capital income, and/or social transfer programs may have the effect of increasing rates of absolute mobility.

Family Structure.—A second specification decision has to do with how to compare incomes across different family structures. In our baseline analysis, we simply

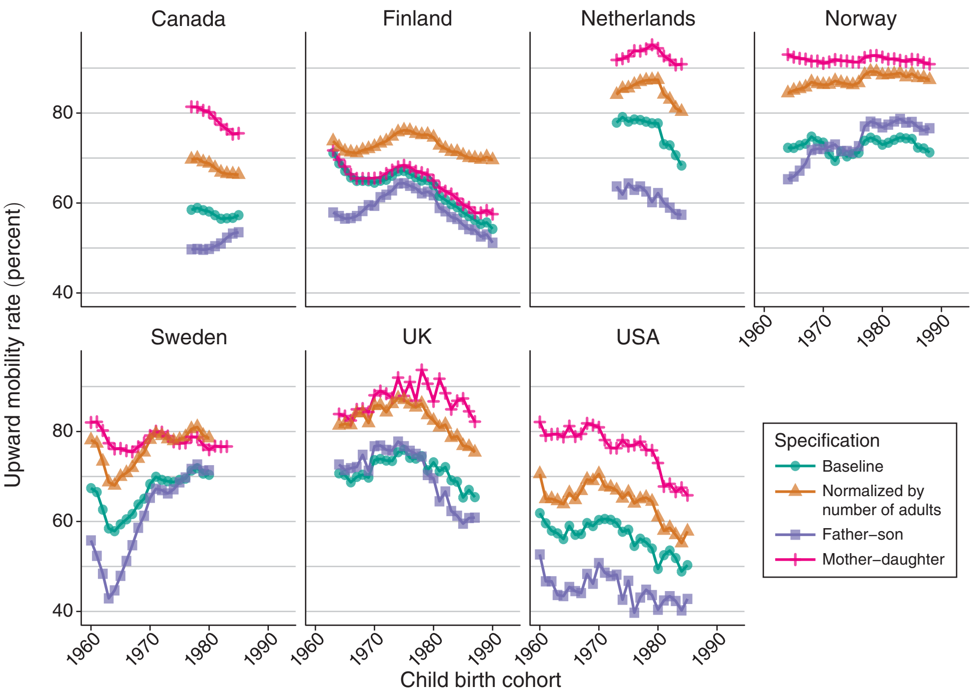

Figure 7. Family Structure Analysis

Note: This figure compares our baseline specification with analyses using alternative family structures.

sum the incomes of spouses to construct a measure of total family income. This may be the most salient income figure in an individual's mind, since it is the amount coming into their household every year. However, this approach doesn't take into account how differences in family structure will impact both the total amount of family income and the standard of living that that income level allows (Bloome 2014).

We conduct three alternative specifications to explore the impact of family composition and labor force participation by gender on our results, shown in Figure 7. First, we analyze mobility after normalizing income by the number of adults in each family. In practice this means dividing total income by two for couples while leaving income for singles unchanged. This specification accounts for the possibility that differences in total income may be due to different family structures rather than differences in the earnings patterns of individuals. There has been a secular decline in marriage among countries in our sample over the past several decades (Lesthaeghe 2010; Ruggles 2015). If members of younger cohorts are remaining single at higher rates than those of older cohorts, their total family incomes may be lower simply because there are fewer adults in the household (cf. Bloome 2014; Western, Bloome, and Percheski 2008). This specification also accounts for differences in family structure between children and parents that are induced by the fact that we include all members of child cohorts in our analysis but only those members of the parent cohorts who had children, who are more likely to be coupled. As shown in the summary statistics in online Appendix Table A3.1, in most sample countries a larger proportion of parents than children were married or cohabiting.

The orange triangles in Figure 7 show that across countries and cohorts in our sample, upward mobility using the normalized income measure is consistently 8–17 percentage points higher than baseline. This suggests that different family structures have resulted in lower family incomes for our sample children than would pertain if they had family structures similar to their parents.

The other two specifications we consider compare the individual incomes of fathers and sons and mothers and daughters, respectively, rather than total family income. This comparison isolates the mobility patterns that are due to earnings trends among men or women alone from those that are due to changes in overall economic conditions or family structure.

In Canada, Finland, Norway, Sweden, and the United Kingdom, upward mobility rates for sons compared to fathers are similar to those using total family income, although the dips in mobility associated with the early 1990s recession in Sweden and the global financial crisis in the United Kingdom are steeper when looking at father–son mobility. In the Netherlands and the United States, upward mobility rates for sons alone are consistently 10–20 percentage points lower than those using total family income. In the Netherlands, this likely reflects the massive increase in female labor force participation since the 1970s, by far the largest in the OECD (Olivetti and Petrongolo 2017). In the United States, it may reflect the faster-than-average decline in male labor force participation over this same period (Krause and Sawhill 2017) or the partial closing of the gender earnings gap (Blau and Kahn 2017).

In all countries except Finland, upward mobility rates for daughters compared to mothers are substantially higher than those using total family income, likely reflecting increasing labor force participation among women over the course of the late twentieth century. The largest recent gaps appear in the Netherlands, Norway, and the United Kingdom, and the smallest in Sweden. In Norway, mobility rates in individual income for both sons and daughters are higher than those in family income. This may result from the combination of Norway's consistently low marriage rate among children and the fact that due to the method of data construction, Norwegian children in our sample are matched to both biological parents even if they only live with one of them (see online Appendix 2). That would be consistent with the large gap between the baseline and adults-normalized family income series in Norway as well.

Price Index.—The final specification choice that we consider regards the price index used to adjust for inflation. The proper way to measure inflation has long been debated (Abraham, Greenlees, and Moulton 1998; Boskin et al. 1997). Challenges include how to account for consumers' substitution of goods due to changes in price, how to quantify the benefits of technological advances, and how to create one summary index that is valid for people with a range of income levels and purchasing habits (Jaravel 2019).

In our baseline analysis, we use the official CPI for each sample country. While a perfect price index is impossible, governments must construct some measure of

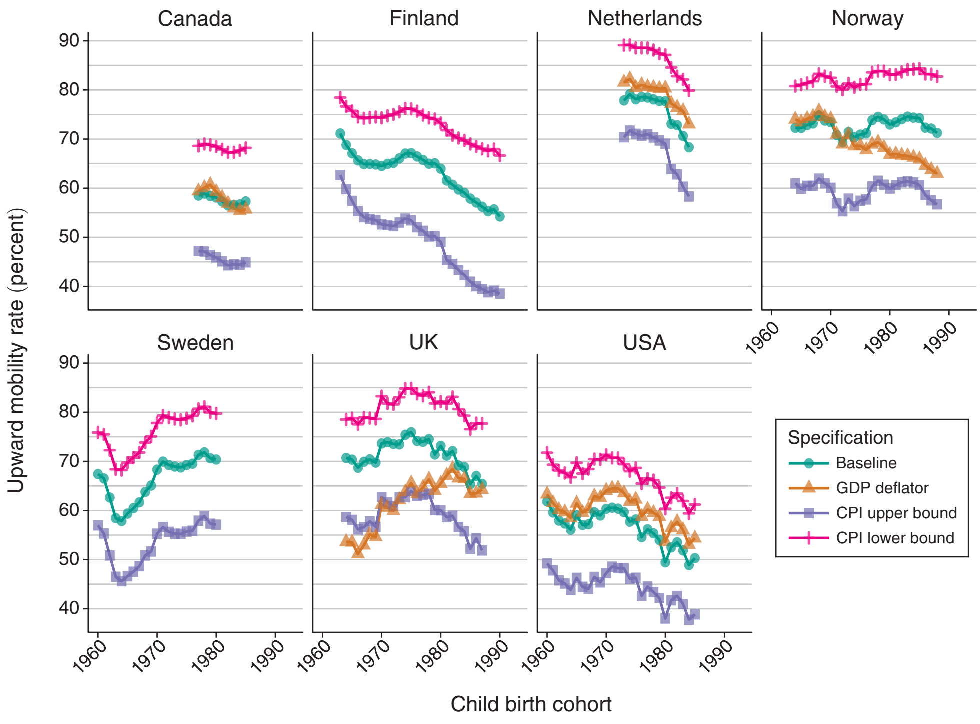

Figure 8. Alternative Price Indices

Notes: This figure compares our baseline specification to specifications using alternative price indices. GDP deflators covering the full sample period are not available for Finland or Sweden. CPI upper bound adds 1 percentage point per year to the inflation rate calculated with the CPI. CPI lower bound subtracts 1 percentage point per year from the inflation rate measured with the CPI.

changes in the price level for the purposes of macroeconomic policymaking, and we use that measure. However, there are often concerns about the official price index in particular countries, and US studies (Chetty et al. 2017; Winship 2017) have shown that estimated absolute mobility rates are sensitive to the price index used.

In this analysis, we explore the sensitivity of upward mobility rates to three alternative price indices. The most widely available alternative to the CPI is the GDP deflator, which weights changes in prices by share of GDP rather than a particular consumption basket. For most conceptions of absolute mobility, personal consumption rather than GDP is likely to be the more relevant weighting mechanism, but the GDP deflator may nonetheless be a useful robustness check.

Of the countries in our sample, five have GDP deflators covering the full sample period. Absolute mobility trends using the GDP deflator are plotted in the orange triangles in Figure 8. The impact of using this deflator varies from country to country: In the United States, Canada, and the Netherlands, using the GDP deflator does not substantially alter trends, while shifting the level of absolute mobility up a few percentage points. In the United Kingdom, absolute mobility using the GDP deflator is substantially lower for early cohorts but rises to be comparable to the baseline series for recent birth cohorts. In Norway, the opposite pattern occurs: upward mobility using the GDP deflator shows a steady downward trend for cohorts born after 1970. For the Norwegian results in particular, we suggest extreme caution when interpreting the GDP deflator. Norway is a small, open economy that is heavily dependent on oil, which accounted for 14 percent of Norwegian GDP in 2021 (BBC 2021). Changes in the price of oil will thus have a substantial impact on the Norwegian GDP deflator while not necessarily translating into similar changes in the prices facing Norwegian consumers.

We also construct two synthetic CPI series in each country that add and subtract 1 percentage point to the inflation rate using the baseline CPI in each year. We believe that these series serve as plausible upper and lower bounds on the amount of inflation that might have occurred between parent and child income measurement. For all countries, trends in absolute mobility look similar in the "CPI plus or minus 1 percentage point series" to how they look in the baseline, but are between 10 and 15 percentage points lower and higher, respectively.

Measuring inflation accurately is critical to the estimation of absolute income mobility, but poses both conceptual and empirical challenges—conceptual because there is no perfect price index, and empirical because it is difficult to construct a price index from scratch and in any given country there may be only a few official price indices available. In general, we recommend using a price index tied to consumer rather than producer prices, since those are likely to be more salient for individuals. In many cross-national applications, it may make sense to use the official CPI, since any errors induced by using that index will also occur in other comparative research. But in individual countries where well understood alternatives to the official CPI are available, researchers may wish to use those instead.

In this section we have explored the sensitivity of absolute mobility rates in all seven sample countries to a range of specification choices relating to the age and length of income measurement, income concept, family structure, and price index. Interestingly, in many cases the impact of a particular specification differs from country to country. Because of this, we recommend care both when comparing absolute mobility rates from one country to another and when choosing a baseline specification, even in studies of a single country.

VI. Decomposition of Sources of Variation in Absolute Mobility across Countries and Cohorts

Thus far, we have documented substantial variation in rates of upward absolute income mobility, both between countries and over time. More than three-quarters of children born in the mid-1980s in Norway, for example, grew up to earn more than their parents at age 30, compared to just half of US children born during this time—and less than 60 percent of children in Finland, next door. An important question for researchers and policymakers is what accounts for this variation. In this final methodological section, we describe a framework for decomposing the differences in upward mobility rates between any two cohorts—whether children born in the same year in different countries or those born in the same country in different years—and discuss some opportunities and challenges in its interpretation. This method can also be used to determine the source of differences between specifications: is the reason that US absolute mobility in disposable income is higher than that in pretax income because the former displays less inequality among children or higher average growth rates?

The absolute mobility rate of a given cohort in a given country can be fully accounted for by four components: the copula, the ratio of mean income in the child generation to mean income in the parent generation, and the shape of the income distribution (that is, the level of inequality) in (i) the parent and (ii) the child generations (Liss et al. 2023; Van Kerm 2004). Using the "copula and marginals" approach, it is possible to determine the fraction of the difference in absolute mobility between two cohorts of children due to each of these components by conducting counterfactual simulations in which each component is varied individually.

A few caveats apply to this method, however. First, while the substitution of all four components for one cohort with those for another cohort will result in perfectly matching the latter's rate of absolute mobility, the sum of the changes induced by each individual substitution will not equal the overall difference in absolute mobility. Second, if substitutions are applied sequentially, the order of application will influence the amount of the difference that appears to be due to a particular component. For that reason we generally recommend against a sequential substitution and in favor of individual substitutions of one or a combination of components.

Figure 9 illustrates the possible uses of the decomposition by decomposing the difference in upward absolute income mobility between the 1985 US birth cohort in our baseline specification and four alternative cohorts: the most recent Norwegian and Swedish cohorts, the 1960 US birth cohort, and the 1985 US birth cohort measured using disposable income. As the figure shows, in all four cases a majority of the gap is closed by substituting the child-to-parent mean income ratio from the higher mobility cohort for that from the 1985 US cohort. Unfortunately, this is the component that is influenced by the largest number of independent factors. The growth of mean incomes from the parent to the child generation is influenced by the overall economic growth rate, but also by the distribution of income across ages, the fraction of GDP accounted for in our income measures, and potentially by changes in typical family structures from parents to children. This is also where measurement error in the form of underreported income is likely to come in. Perhaps less immediately apparent, it is also influenced by the average age at which parents had children, since that determines the number of years of economic growth between when the parents and the children turn 30.

The second-most important component, particularly in the US disposable income comparison, is the amount of inequality among the child cohort. Substituting the inequality among parents actually exacerbates the gap, likely because lower parent inequality results in higher parent incomes across much of the distribution, which will be associated with lower upward mobility, all else equal. The copula has minimal impact, though this is by default in the 1960 US and 1985 US disposable income scenarios since those use the same copula as the 1985 US baseline. This is consistent with our findings reported in Section IV that variation among empirical copulas has relatively limited impact on the rate of absolute mobility.

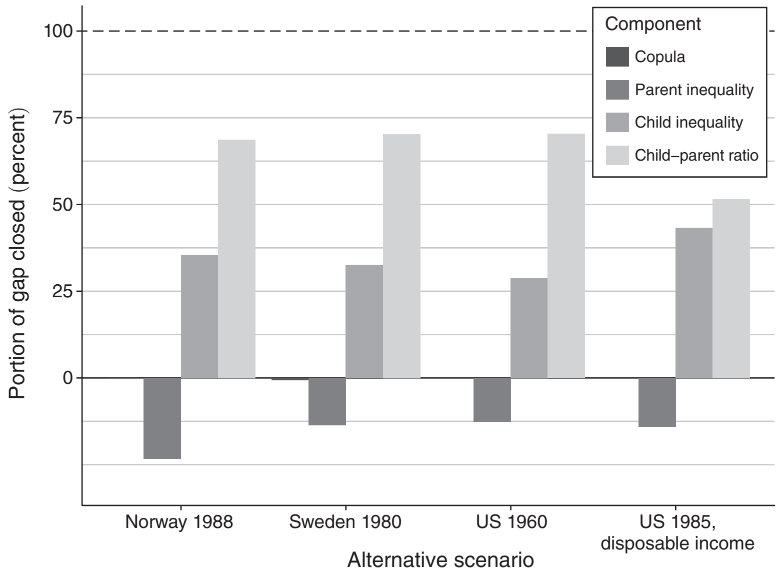

Figure 9. Decomposition of Differences in Absolute Mobility Between the US 1985 Birth Cohort and Various Alternative Scenarios

Notes: This figure decomposes the sources of the difference in upward mobility between the baseline US 1985 birth cohort (upward mobility rate of 50.3 percent) and 4 other cohorts: the Norway 1988 cohort (71.3 percent), the Sweden 1980 cohort (70.4 percent), the US 1960 cohort (61.9 percent), and the US 1985 cohort measured using disposable income (59.3 percent). In each case, the bars show the percentage of the overall gap closed by substituting the copula, parent inequality, child inequality, and child–parent mean income ratio, respectively, from the high-mobility cohort for that in the 1985 US baseline cohort. When substituted cumulatively, these four components perfectly account for the entire difference in upward mobility.

The patterns in Figure 9 appear to be typical of the countries and cohorts in our sample. As a supplemental exercise, shown in online Appendix Figure A8.1, we conducted this decomposition for all pairs of cohorts in our baseline sample that differed in their upward mobility rates by more than 10 percentage points. Across these 2,008 pairs of cohorts, the median percentage of the gap attributable to the child–parent mean income ratio was 66 percent, that attributable to inequality among children was 37 percent, that attributable to inequality among parents was –11 percent, and that attributable to the copula was –2 percent.

An important task for research on absolute income mobility is to determine why mobility rates rise and fall. Decomposition methods like those outlined here are one major way in which this can be done. However, their interpretation is not always straightforward, so we encourage researchers to exercise care when conducting decompositions of this sort.

VII. Discussion

In this paper we have explored a range of methodological issues in the calculation and interpretation of absolute income mobility. Using linked parent–child administrative data for five countries in North America and Europe, we demonstrated that the increasingly common "copula and marginals" approach to estimating rates of absolute mobility without linked data closely tracks the results calculated directly from linked records for the countries and cohorts where we have access to such records. We also presented evidence that the assumption of copula stability—used to justify applying the copula from a recent birth cohort to marginal distributions from earlier cohorts—is unlikely to cause substantial error. More troubling is the possibility of error introduced when creating discrete 100- or (especially) 10-cell marginal distributions for children and parents. The largest and least predictable source of error, however, is likely to come from the use of alternate populations to create the marginal distributions. For this reason we recommend that researchers do as much as possible to validate that their income data closely match those for the target populations—meaning populations of the same age and family status. While even more straightforward methods of approximating absolute mobility based on median incomes alone can be surprisingly accurate (Katz and Krueger 2017), we recommend caution when using them, especially when comparing to absolute mobility rates constructed from linked records or the "copula and marginals" method.

We have also performed a multicountry analysis of the sensitivity of absolute mobility to a range of specification choices. Among the cohorts in our sample, absolute mobility rates appear to stabilize when income is measured after age 30 in certain countries and after age 35 in all countries. In most of the countries we were able to test, absolute mobility in income averaged over multiple years is slightly higher than, but within a few percentage points of, mobility calculated from a single year of income—but with some indications that multiyear incomes are less affected by business cycle fluctuations.

Other specifications that we considered showed variation in their effects by country in ways that highlight differences in political economy and demographic trends. Controlling for family size had the largest impact in Norway, where the parent–child difference in couple rates is greatest (likely due in part to the method of linking parents and children, as described in online Appendix 2), while the difference between mobility in family income and that in individual income for both sons and daughters was greatest (in opposite directions) for the Netherlands, which has seen a particularly marked increase in female labor force participation since the 1970s.

Finally, we offered some guidance on the use of decomposition methods to determine the sources of mobility differences between countries, cohorts, and even specifications. Unfortunately, the most analytically straightforward method is not well suited to distinguishing between many of the most likely possible sources of variation in mobility rates, which differ in their substantive interpretations and implications for policy. But distinguishing between within-cohort child and parent inequality and income growth from parents to children does provide important insight in certain cases.

Absolute income mobility remains a highly salient and revealing economic indicator, one that deserves continued attention from social scientists and policymakers. It is our hope that the tools we have worked to validate will help future researchers to accurately portray and analyze the levels and determinants of upward absolute income mobility, so that future generations of children may continue to realize the dream of a higher standard of living than their parents.

REFERENCES

Abraham, Katharine G., John S. Greenlees, and Brent R. Moulton. 1998. "Working to Improve the Consumer Price Index." *Journal of Economic Perspectives* 12 (1): 27–36.

Amiel, Yoram, and Frank A. Cowell. 1999. *Thinking about Inequality: Personal Judgment and Income Distributions.* Cambridge, UK: Cambridge University Press.

BBC. 2021. "Norway's Oil and Gas Sector Will Not Be Dismantled, New Government Says." *BBC News,* October 13. https://www.bbc.com/news/world-europe-58896850.

Berman, Yonatan. 2022. "The Long Run Evolution of Absolute Intergenerational Mobility." *American Economic Journal: Applied Economics* 14 (3): 61–83.

Blau, Francine D., and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55 (3): 789–865.

Bloome, Deirdre. 2014. "Racial Inequality Trends and the Intergenerational Persistence of Income and Family Structure." *American Sociological Review* 79 (6): 1196–1225.

Bönke, Timm, Astrid Harnack, and Holger Luthen. 2019. "Are Children Better Off? Intergenerational Mobility of Living Standards." Unpublished.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson. 1997. "The CPI Commission: Findings and Recommendations." *American Economic Review Papers and Proceedings* 87 (2): 78–83.

Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel D. Schnitzlein, and Kjell Vaage. 2017. "A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US." *Scandinavian Journal of Economics* 119 (1): 72–101.

Bukodi, Erzsébet, Marii Paskov, and Brian Nolan. 2020. "Intergenerational Class Mobility in Europe: A New Account." *Social Forces* 98 (3): 941–72.

Bynner, J., N. Butler, and University of London, Institute of Education, Centre for Longitudinal Studies. 2022. "1970 British Cohort Study: Age 16, Sweep 4, 1986" (accessed January 16, 2023).

Canberra Group. 2011. *Handbook on Household Income Statistics.* Geneva: United Nations.

Central Statistical Office. 1997. "Family Expenditure Survey, 1989 to 1995–1996" (accessed January 16, 2023).

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *Quarterly Journal of Economics* 129 (4): 1553–1623.

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 Papers and Proceedings* 104 (5): 141–47.

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.

Chuard, Patrick, and Veronica Grassi. 2020. "Switzer-Land of Opportunity: Intergenerational Income Mobility in the Land of Vocational Education." Unpublished.

Corak, Miles. 2016. "Inequality from Generation to Generation: The United States in Comparison." IZA Discussion Paper 9929.

Department for Work and Pensions, Office for National Statistics, NatCen Social Research. 2018. "Family Resources Survey, 1993–1994 to 2018–2019" (accessed January 16, 2023).

Department of Employment. 1990. *Family Expenditure Survey, 1968–1988.* London, UK: HMSO.

Deutscher, Nathan, and Bhashkar Mazumder. 2023. "Measuring Intergenerational Income Mobility: A Synthesis of Approaches." *Journal of Economic Literature* 61 (3): 988–1036.

Engzell, Per, and Carina Mood. 2021. "How Robust Are Estimates of Intergenerational Income Mobility?" Unpublished.

Erikson, Robert, and John H. Goldthorpe. 1992. *The Constant Flux: A Study of Class Mobility in Industrial Societies.* Oxford, UK: Oxford University Press.

Eshaghnia, Sadegh, James Heckman, Rasmus Landersø, and Rafeh Qureshi. 2022. "Intergenerational Transmission of Family Influence." Unpublished.

Esping-Andersen, Gøsta. 1990. *The Three Worlds of Welfare Capitalism.* Princeton, NJ: Princeton University Press.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westburry. 2022. "Integrated Public Use Microdata Series, Current Population Survey" (accessed July 22, 2022).

Grawe, Nathan D. 2006. "Lifecycle Bias in Estimates of Intergenerational Earnings Persistence." *Labour Economics* 13 (5): 551–70.

Harding, David J., and Martin D. Munk. 2020. "The Decline of Intergenerational Income Mobility in Denmark: Returns to Education, Demographic Change, and Labor Market Experience." *Social Forces* 98 (4): 1436–64.

Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. "Who Suffers During Recessions?" *Journal of Economic Perspectives* 26 (3): 27–48.

Jäntti, Markus, and Stephen P. Jenkins. 2015. "Chapter 10—Income Mobility." In *Handbook of Income Distribution,* Vol. 2, edited by François Bourguignon and Anthony B. Atkinson, 807–935. Amsterdam: Elsevier.

Jaravel, Xavier. 2019. "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector." *Quarterly Journal of Economics* 134 (2): 715–83.

Katz, Lawrence F., and Alan B. Krueger. 2017. "Documenting Decline in US Economic Mobility." *Science* 356 (6336): 382–83.

Kennedy, Tomas, and Peter Siminski. 2022. "Are We Richer than Our Parents Were? Absolute Income Mobility in Australia." *Economic Record* 98 (320): 22–41.

Krause, Eleanor, and Isabel V. Sawhill. 2017. *What We Know and Don't Know about Declining Labor Force Participation: A Review.* Washington, DC: Brookings Institution.

Lesthaeghe, Ron. 2010. "The Unfolding Story of the Second Demographic Transition." *Population and Development Review* 36 (2): 211–51.

Liss, Erik, Martin Korpi, and Karl Wennberg. 2023. "Absolute Income Mobility and the Effect of Parent Generation Inequality: An Extended Decomposition Approach." *European Economic Review* 152: https://doi.org/10.1016/j.euroecorev.2022.104359.

Luxembourg Income Study. 2023. "Luxembourg Income Study (LIS) Database" (accessed July 20, 2022).

Manduca, Robert, Maximilian Hell, Adrian Adermon, Jo Blanden, Espen Bratberg, Anne C. Gielen, Hans Van Kippersluis, et al. 2024. "Replication Data for: Measuring Absolute Income Mobility: Lessons from North America and Europe." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E176522V1.

Mazumder, Bhashkar. 2005. "Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data." *Review of Economics and Statistics* 87 (2): 235–55.

Nybom, Martin, and Jan Stuhler. 2016. "Heterogeneous Income Profiles and Lifecycle Bias in Intergenerational Mobility Estimation." *Journal of Human Resources* 51 (1): 239–68.

Nybom, Martin, and Jan Stuhler. 2017. "Biases in Standard Measures of Intergenerational Income Dependence." *Journal of Human Resources* 52 (3): 800–25.

Office for National Statistics. 2002. *Family Expenditure Survey, 1996–1997 to 2000–2001.* London, UK: HMSO.

Olivetti, Claudia, and Barbara Petrongolo. 2017. "The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries." *Journal of Economic Perspectives* 31 (1): 205–30.

Ostrovsky, Yuri. 2017. "Doing as Well as One's Parents? Tracking Recent Changes in Absolute Income Mobility in Canada." Statistics Canada. https://www150.statcan.gc.ca/n1/pub/11-626-x/11-626-x2017073-eng.htm (accessed September 5, 2022).

Pekkarinen, Tuomas, Kjell G. Salvanes, and Matti Sarvimäki. 2017. "The Evolution of Social Mobility: Norway during the Twentieth Century." *Scandinavian Journal of Economics* 119 (1): 5–33.

Ravallion, Martin. 2018. "Inequality and Globalization: A Review Essay." *Journal of Economic Literature* 56 (2): 620–42.

Ruggles, Steven. 2015. "Patriarchy, Power, and Pay: The Transformation of American Families, 1800–2015." *Demography* 52 (6): 1797–1823.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek. 2022. "IPUMS USA: Version 12.0." IPUMS (accessed May 10, 2022).

Sirniö, Outi, Timo M. Kauppinen, and Pekka Martikainen. 2017. "Cohort Differences in Intergenerational Income Transmission in Finland." *Acta Sociologica* 60 (1): 21–39.

Sklar, Abe. 1959. "Fonctions de Repartition an Dimensions et Leurs Marges." *Publications de l'Institut Statistique de l'Université de Paris 8*: 229–31.

Smeeding, Timothy M., Robert Erikson, and Markus Jäntti. 2011. *Persistence, Privilege, and Parenting: The Comparative Study of Intergenerational Mobility.* New York, NY: Russell Sage Foundation.

Solon, Gary. 1992. "Intergenerational Income Mobility in the United States." *American Economic Review* 82 (3): 393–408.

Statistics Canada. 2019. "Intergenerational Incomes Database (IID)." Statistics Canada (accessed June 5, 2022).

Statistics Finland. 2020. "Products and Services." Statistics Finland. http://www.stat.fi/tup/index_en.html (accessed July 1, 2020).

Statistics Netherlands. 2020. "Microdata: Conducting Your Own Research." Statistics Netherlands. https://www.cbs.nl/en-gb/our-services/customised-services-microdata/microdata-conducting-your-own-research (accessed July 1, 2020).

Statistics Norway. 2018. "Data for Research." Statistics Norway. https://www.ssb.no/en/data-til-forskning (accessed July 1, 2020).

Statistics Sweden. 1800–2016. "Befolkningsregistret (RTB) (Population Register) IFAU-U Version." IFAU (accessed October 19, 2017).

Statistics Sweden. 1932–2016. "Flergenerationsregistret (Multigenerational Register), IFAU-U Version." IFAU (accessed October 19, 2017).

Stockhausen, Maximilian. 2018. "Like Father, Like Son?—A Comparison of Absolute and Relative Intergenerational Labour Income Mobility in Germany and the US." *Journal of Economic Inequality* 19: 667–83.

Torche, Florencia. 2015. "Analyses of Intergenerational Mobility: An Interdisciplinary Review." *ANNALS of the American Academy of Political and Social Science* 657 (1): 37–62.

University of London. 2013. *1970 British Cohort Study: Twenty-Nine-Year Follow-Up, 1999–2000.* London, UK: Centre for Longitudinal Studies.

van Kerm, Philippe. 2004. "What Lies behind Income Mobility? Reranking and Distributional Change in Belgium, Western Germany, and the USA." *Economica* 71 (282): 223–39.

Western, Bruce, Deirdre Bloome, and Christine Percheski. 2008. "Inequality among American Families with Children, 1975 to 2005." *American Sociological Review* 73 (6): 903–20.

Winship, Scott. 2017. *Economic Mobility in America: Contemporary Levels of Mobility.* Washington, DC: Archbridge Institute.