

Who Gets Ahead? Measuring Income Gaps across Family Backgrounds and Neighborhoods

Sander de Vries*

June 21, 2026

Abstract

This paper introduces new approaches to measure intergenerational transmission with detailed family information. I link 1.7 million Dutch children to 91 family-background characteristics to provide the first estimates of how long-run income varies across the full family-background distribution, including the tails. Compared to traditional approaches using parental income only, I document considerably larger income gaps across families and smaller gaps across neighborhoods. The mean-income gap between children from the least and most advantaged families widens from 42 to 61 ranks, while cross-neighborhood variation in upward mobility falls by over half. These results matter for policymakers targeting disadvantaged families and neighborhoods.

Keywords: intergenerational mobility, inequality of opportunity

JEL Codes: J24, J62

*Department of Economics, Vrije Universiteit Amsterdam, s.de.vries@vu.nl. I gratefully acknowledge valuable comments from Nadine Ketel, Maarten Lindeboom, Erik Plug, Paul Hufe, Gustave Kenedi, Stanislav Avdeev, and conference and seminar participants in Amsterdam, Utrecht, The Hague, Tokyo, London, Canazei, Colchester, Antwerp, Sydney, and Groningen. Earlier versions of this paper were circulated under the title ‘Measuring Family (Dis)Advantage: Lessons from Detailed Parental Information’. The non-public micro data used in this paper are available via remote access to the Microdata services of Statistics Netherlands (project agreement 8674).

1 Introduction

The income prospects of the children from the least and most advantaged families are central to debates about equality of opportunity. Low upward mobility is widely seen as unfair and a justification for policy intervention (Roemer and Trannoy (2016), Alesina et al. (2018)). At the top, the perpetuation of advantage raises equally important concerns about whether economic success reflects merit or inheritance. Yet, despite its importance, there is still a lot we do not know about how income is distributed across family backgrounds.

Recent administrative data enabled researchers to describe in detail the income prospects of the children growing up in the lowest- and highest-income families (e.g., Chetty et al. (2014), Deutscher and Mazumder (2020), Kenedi and Sirugue (2023)). However, even among low-income families, some children face especially poor prospects when the family is also disadvantaged in other dimensions such as education or family structure. Univariate analyses mask the low income of these children. While measures such as sibling correlations or inequality of opportunity estimates aim to capture broader family influences, they summarize income disparities across the family background distribution in a single population-level statistic. No existing study has offered an analysis that is both *broad* — incorporating many family dimensions — and *local* — describing how income varies at specific points of the family background distribution, including the tails.¹

This paper fills this gap by addressing two main challenges. The first is the demanding data requirements. Measuring family background broadly requires observing many characteristics beyond parental income, while providing local estimates demands a much larger sample than population-level analysis. I address this by exploiting Dutch administrative data that link the long-run incomes of 1.7 million children to 91 family background variables. These variables contain detailed information on parents and extended family, including income, wealth, occupation, education, crime, healthcare, migration background, and family struc-

¹Deutscher and Mazumder (2023) classify various measures of intergenerational mobility and introduced the terms *broad* and *local*.

ture. Although this dataset does not capture all potentially relevant dimensions of family background, it greatly expands upon the scope of previous studies.

The second challenge is to translate this high-dimensional information into interpretable measures that reveal local insights. As in the recent inequality of opportunity literature (Brunori et al. (2023)), I use all family information in a flexible prediction model of child income. Instead of summarizing the predicted disparities in one inequality-of-opportunity index, I provide the full distribution of expected income and probabilities of reaching different income quintiles. These are the multidimensional analogues of the conditional expectation functions and transition matrices common in the intergenerational mobility literature.

This approach reveals strikingly large income gaps between the most and least advantaged children. Among the 0.5 percent of children with the lowest expected income, the average income rank is 17 (out of 100). Their chance of reaching the top income quintile is below one percent. By contrast, the average income rank of the top 0.5 percent of children is 78, with 64 percent reaching the top income quintile. This gap of 61 ranks amounts to a gross household income difference of 146,000 euros annually. An analysis with only parental income strongly underestimates this gap: ranking children by just parental income, the bottom and top 0.5 percent of children differ by only 42 ranks (or 54,000 euros annually).

I provide detailed insights into which family characteristics most strongly predict child income using a child-level Shapley value decomposition (Lundberg et al. (2020)). Conventional decompositions typically quantify how much of the population-wide variation in income each variable explains (Shorrocks (2013)).² The child-level Shapley decomposition instead quantifies how much of each child's expected income a variable explains. To my knowledge, this is the first application of Shapley decompositions at the individual level to study income inequality. It offers two advantages. First, it shows for how many children, and by how much, each family characteristic contributes to expected income. Second, it enables decomposing the expected income for specific subgroups, such as the most disadvantaged children.

²Some of these approaches also use Shapley value decompositions, for example to quantify variable contributions to R^2 in linear regression or to decompose inequality (of opportunity) indices.

Income and wealth of parents and extended family make the largest contributions on average. Together they explain 78 percent of the variation in expected income, and are therefore most important for quantifying family-driven inequality at the population level. At the local level, among the most disadvantaged children, family income and wealth explain a relatively smaller share. The low expected income of these children reflects the cumulative negative contribution of unfavorable family characteristics across many dimensions: beyond low income and wealth, their parents are also often young, separated, minimally educated, suspected of crimes, and with high health expenditures and a migration background.

Finally, I study how income prospects vary across fine-grained neighborhoods. A standard measure of neighborhood upward mobility based on parental income alone varies considerably across neighborhoods. However, I show that much of this variation reflects sorting on family characteristics other than income. To correct for this, I introduce a new measure of neighborhood upward mobility using all family information. Cross-neighborhood variation in expected income falls by 55 percent under this measure, and the ranking of neighborhoods changes substantially: neighborhoods with similar rankings under the income-based measure differ by 20 ranks on average (out of 100) under the multidimensional measure. This matters because sorting alone would not distort policy targeting if the ranking of high- and low-mobility neighborhoods were preserved.

This paper contributes to the growing literature measuring intergenerational transmission with multiple observable family background characteristics. New approaches measuring intergenerational persistence (Vosters and Nybom (2017), Vosters (2018), Adermon et al. (2021), Eshaghnia et al. (2022), Chang et al. (2025)), inequality of opportunity (Brunori et al. (2023), Adermon et al. (2025)), or explanatory power (Blundell and Risa (2019), Althoff et al. (2025)) all document how family background shapes children’s economic success. Some papers also decompose the contributions of family background characteristics to these measures (Blundell and Risa (2019), Salas-Rojo and Rodríguez (2022), Brunori et al. (2023), Althoff et al. (2025)). The approaches above are broad but *global*: they capture multiple

family background dimensions but focus on population-level statistics, such as overall persistence or total variance explained by each variable.³ The main contribution of this paper is to present an analysis that is broad and *local*: I provide the first estimates of expected income — and decompose the contribution of each family characteristic to it — across the full distribution of family backgrounds, with particular focus on children at the tails.⁴

Local insights matter because fairness concerns depend not only on average persistence, but also on whether some children face particularly low chances of upward mobility or inherit exceptional advantages (Roemer and Trannoy (2016)). The results in this paper show that this is the case for a small group of children whose parents are (dis)advantaged in nearly every measurable dimension. The results also have a direct policy implication: interventions that target families on a single characteristic, such as low income, both include children who are not particularly disadvantaged on other dimensions and miss children who are heavily disadvantaged on many dimensions but not on income. The results show that combining several family characteristics substantially improves identification of the most disadvantaged.

A second key contribution is to the literature on spatial differences in income prospects. A growing literature documents large geographic differences in the expected income of children from *low-income* families (e.g., Chetty et al. (2014); Heidrich (2017); Deutscher and Mazumder (2020); Corak (2020); Alesina et al. (2021); Acciari et al. (2022); Kenedi and Sirugue (2023)). My multidimensional measure of neighborhood upward mobility extends Chetty et al. (2026), who estimate neighborhood mobility from income predictions based on the national parent-child income relationship. I generalize their approach by using predictions based on the full set of family characteristics.⁵ The resulting estimates remain easy

³A similar drawback applies to papers that estimate intergenerational persistence indirectly through correlations with relatives, such as siblings or more distant family members (Solon (1999), Bingley and Cappellari (2019), Collado et al. (2023)). They capture broad family influences, including unobserved ones, but summarize their influence into a single share of explained variation.

⁴Most closely related are Brunori et al. (2023) and Brunori et al. (2024), who estimate how expected income varies across several family types. However, their goal is to provide a population-level estimate of inequality of opportunity. Because their survey data contains far fewer background variables and observations, they are unable to provide local estimates for children from the least and most advantaged families.

⁵This is also related to Cholli et al. (2024), who generalize a standard intergenerational mobility model by including observable family and community characteristics and additionally controlling for unobservables

to interpret and move closer to an ideal measure that separates neighborhood effects from family sorting. Policymakers can use them to identify neighborhoods where children with similar observable family backgrounds face persistently lower income prospects.

This paper proceeds as follows. Sections 2 and 3 present the data and the methodology. Section 4 measures income gaps across family backgrounds and conducts the decomposition exercise. Section 5 measures income gaps across neighborhoods. Section 6 concludes.

2 Data

Core analysis sample. I use administrative data from Statistics Netherlands covering the full Dutch population.⁶ The main sample consists of all children born in the Netherlands between 1980 and 1989, excluding 3.3 percent with missing income data, resulting in 1,705,532 observations.

The main outcome in this paper is a child’s long-run gross household income rank. I focus on household income because it provides a reliable measure of economic resources even in the case of non-participation in the labor market and it is commonly used in other intergenerational mobility studies (Chadwick and Solon (2002)). Nevertheless, I also present results using personal income ranks to abstract away from household formation considerations. Household income is observed between 2001 and 2024 and includes income from employment, entrepreneurship, capital, income insurance payments, social security payments, inter-household income transfers (such as alimony), and contributions to social insurance made by both employers and employees.⁷ Income is measured in 2024 euros, adjusting for inflation using the consumer price index.

I construct a proxy for children’s lifetime household income by averaging their household

using a control-function approach. I discuss conceptual and empirical similarities and differences with their approach in Section 5.

⁶Access is granted through a secure remote facility under a confidentiality agreement.

⁷Some children still live with their parents when I measure their income. In these cases, I define the income of the children as their gross personal income and that of the parents as the household income minus the total gross personal income of the children who still live at home.

income from age 30 onward.⁸ This approach reduces measurement error from transitory income shocks Mazumder (2005) and life-cycle bias (Haider and Solon (2006), Nybom and Stuhler (2017)). I observe income up to age 44 for the oldest cohort (born in 1980) and up to age 35 for the youngest cohort (born in 1989). On average, children have ten income observations, with 96 percent having at least five. I then rank children within birth-years based on their lifetime household income. I also present results for various alternative measures to evaluate the sensitivity of the results due to these choices.

Parental household income. The parent-child register enables me to link children to their legal parents. I then estimate each parent’s lifetime household income by averaging their annual household incomes after 2001 and up to age 60. Since most parents were born in the 1950s, their first incomes are typically observed around their late 40s. On average, fathers have 13 income observations and mothers 16. Following Chetty et al. (2014), parental income is defined as the average of the father’s and mother’s lifetime household income. If only one parent’s income is observed, I use that parent’s income. The parental income rank is based on the position within the parental income distribution of all children in the analysis sample.

Other explanatory variables. Table 1 describes how the other variables are classified into nine categories. Except for household income and wealth, which are measured at the household level, all variables are included for the father and the mother separately. Altogether, the set comprises 75 continuous variables, 8 binary indicators, and 8 categorical variables (two containing 68 distinct categories and six containing 8 categories). Appendix A provides descriptive statistics for the core sample, including all explanatory variables, as well as a detailed explanation of how the explanatory variables are constructed.

Although the data are rich, they come with two limitations. First, some parental outcomes are observed only after their children have left the household. Consequently, my

⁸I exclude years with yearly household income below €1,000 (0.6%), as these cases typically correspond to wealthy entrepreneurs with business losses.

Table 1: Explanatory Variables

Income	Household income, personal income, personal earnings, most important sources of personal income (in 11 categories), and the primary household income share.
Wealth	The value of bank and savings balances, bonds and shares, real estate, entrepreneurial assets and liabilities, other assets, mortgage debt, study debt, and other debt.
Occupation	Average hourly wage and most important sector of employment (in 68 categories).
Education	Highest level of completed education.
Healthcare	Average healthcare costs for 5 categories*: general practitioner, hospital, pharmaceutical, mental health care, and dental care.
Crime	Indicators of whether the parent has been suspected of a property, violent, or other type of crime.
Family structure	Parents' family size, age-at-first-birth, birth order, single-parent household, father or mother presence, parental death, child family size, and whether the father or the mother are identified.
Migration background	Region of origin of the father, mother, and all grandparents (in 8 categories).
Extended family outcomes	Average years of education, average household income rank, average wealth rank, average healthcare costs, and share of all siblings of the parent who have been suspected of a crime.

Notes: this Table describes the explanatory variables used in the main analysis. A detailed explanation of each of the variables and descriptive statistics can be found in Appendix A. *: Healthcare costs are based on healthcare insurance reimbursements. Basic healthcare insurance is mandatory for all residents and covers a wide range of medical services (see also Appendix B).

results may underestimate the importance of family background compared to a model that includes information on parents' resources and well-being during their children's formative years. Nonetheless, many parental characteristics are highly persistent over the life-cycle, making them a reasonable proxy for the family environment at earlier ages.⁹

Second, despite the extensive coverage of variables, some missing values persist. Most importantly, education records for the parents' generation are incomplete. In a robustness check, I assess the impact of these missing education records. Extended family outcomes are also unavailable for some children, often because their parents have no siblings or their

⁹This is supported by Eshaghnia et al. (2025), who show that differences in intergenerational mobility estimates due to different types of resources being analyzed are much larger than differences due to the age of the children at which these resources are measured.

grandparents cannot be identified, making it impossible to link to aunts or uncles. To preserve the full sample, I use indicators to denote missing information instead of excluding incomplete observations.

3 Methodology

The goal is to describe how income varies across family backgrounds. Let Y_{sf} be the income rank of a child s in a family f . Moreover, let $\mathbf{X}_f = (X_{f1}, \dots, X_{f91}) \subset \mathcal{X}$ be the set of all observed family background characteristics. Consider the following conditional expectation function decomposition of Y_{sf} :

$$Y_{sf} = E[Y_{sf}|\mathbf{X}_f] + \nu_{sf} = g(\mathbf{X}_f) + \nu_{sf}, \quad (1)$$

where, by construction, $E[\nu_{sf}] = E[\nu_{sf}m(\mathbf{X}_f)] = 0$ for any function $m : \mathcal{X} \rightarrow \mathbb{R}$.

Global measures. The primary objective of this paper is to measure the importance of observable family background characteristics for child income. This is compactly summarized by the share of income variation attributable to differences in $g(\mathbf{X}_f)$ — the conditional mean for individuals with observable family background \mathbf{X}_f — as opposed to residual variation in income ν_{sf} . This corresponds to the nonparametric R^2 of the observables model:

$$R_{y|g}^2 = \frac{V(g(\mathbf{X}_f))}{V(Y_{sf})}. \quad (2)$$

I commonly refer to this metric as the ‘explanatory power’. Deutscher and Mazumder (2023) classify this as a global measure of intergenerational dependence, as it summarizes the importance of family background for the entire population.

The explanatory power is directly comparable to two commonly used alternative global measures of intergenerational dependence: the rank-rank correlation and sibling correlations.

I use this comparison to benchmark the explanatory power of the full model against (i) models based solely on parental income, which provide a lower bound, and (ii) models based on sibling fixed effects, which provide an upper bound for models that include only observable factors shared between siblings (as in this paper).¹⁰

A closely related approach from the inequality of opportunity literature makes similar decompositions as in Equation 2, but typically uses other inequality measures than the variance. This is called the ex-ante approach to quantifying inequality of opportunity.¹¹ This literature treats all background factors beyond an individual’s control as ‘circumstances’. The findings in this paper are specific to inequality of opportunity arising from family circumstances, a subset of all possible circumstances.

Local measures. I present two measures that provide more detailed insight into the size of the income gaps across family backgrounds. I first report the full distribution of expected incomes, $F(X) = P(g(\mathbf{X}_f) \leq X)$, which allows me to identify the expected income of the least and most advantaged children and all groups in between. This is the multidimensional analogue of commonly reported Conditional Expectation Function plots of child income given parental income. Next, I construct detailed matrices that map children with different expected incomes into quantiles of the observed income distribution. This is the multidimensional analogue of commonly reported transition matrices. In the framework of Deutscher and Mazumder (2023), these measures are local because they provide insights for specific subsets of children.

¹⁰This follows because the sibling correlation equals the (adjusted) R^2 of a regression of child income on family fixed effects. Because these family fixed effects measure the importance of *all* factors shared between siblings, including unobserved ones, their explanatory power is necessarily higher than that of any model using only observable factors shared between siblings. The explanatory power of observables can be higher than the sibling correlation if they also include factors that differ between siblings, such as birth order effects, genes, or life-cycle variations in parental income over time.

¹¹A detailed explanation of this and related approaches can be found in Roemer and Trannoy (2016) and Ramos and Van de Gaer (2016). Brunori et al. (2024) also discuss how intergenerational mobility coefficients and inequality of opportunity estimates are related.

Estimation. Estimation of the global and local measures requires estimation of the conditional expectation function $g(\mathbf{X}_f)$. A key challenge is that its functional form is unknown. Variables may enter in a non-linear manner or interact with other variables. In these cases, nonparametric machine learning methods outperform linear regression models. Accordingly, I employ gradient-boosted decision trees to generate these predictions (Friedman (2001)).¹² Tree-based methods offer the additional advantage of providing Shapley value-based measures of variable importance even with a large number of predictors.

The complexity of tree-based models depends on several tuning parameters, such as the maximum number of splits per tree, the minimum gain required for a split, the total number of trees, and the learning rate. To select these parameters, I randomly split the sample into a training set (80 percent) and a test set (20 percent). I use 5-fold cross-validation on the training set to determine optimal values and then re-estimate the model on the full training set with these parameters. I apply the final model to the test set to obtain out-of-sample predictions, from which I compute both global and local measures.

Quantifying variable importance. A next question is how much each of the family background characteristics contribute to differences in expected income between children.

Decomposing family characteristics' contributions to expected income is not trivial. For example, based solely on a predictive model $\hat{g}(\mathbf{X}_f)$ it is not possible to meaningfully quantify variable importance when variables correlate or interact (Hastie et al. (2001)). This is because when variables interact, their joint effect cannot be separated into independent contributions, and when variables correlate, the model can mix up one variable's contribution with that of the other. Considering variables in isolation or removing one variable at a time does not solve this problem either. Variables that appear uninformative on their own may become important when combined with others, while variables that seem redundant in the

¹²Single decision trees partition the covariate space into regions with similar outcomes and predict for new observations the average value in their region. Gradient-boosted trees improve on this by iteratively fitting trees to residuals, enabling them to capture non-linear relationships and complex interactions. Such tree-based methods have proven to be superior when predicting from tabular data (Grinsztajn et al. (2022)).

full model may contribute substantially in smaller subsets.

A solution is to average a variable’s marginal contribution to a prediction over all possible combinations of other variables. The Shapley value, introduced in cooperative game theory, does exactly this (Shapley (1953)). Lundberg and Lee (2017) show that a Shapley decomposition is the only way to quantify variable importance at the individual level while preserving important properties.¹³ While computing such Shapley values is infeasible for most models due to the need to re-estimate models for all possible variable subsets, a recent algorithm can compute them for tree-based models in short time periods (Lundberg et al. (2020)).¹⁴ I use this algorithm to compute Shapley values from the gradient-boosted decision tree. Below, I briefly explain the intuition behind this approach.

First define the marginal contribution of variable X_{fk} to a given subset of variables $S \subseteq \mathbf{X}_f \setminus \{X_{fk}\}$ by the change in expected income induced by adding this variable:

$$h(X_{fk}, S) = E[Y_{sf}|S \cup X_{fk}] - E[Y_{sf}|S].$$

For example, when there are no other explanatory variables ($S = \emptyset$), then the contribution of X_{fk} equals $h(X_{fk}, \emptyset) = E[Y_{sf}|X_{fk}] - E[Y_{sf}]$. At the other extreme, when S is the set of all other variables, then the contribution of X_{fk} equals $h(X_{fk}, \mathbf{X}_f \setminus \{X_{fk}\}) = g(\mathbf{X}_f) - E[Y_{sf}|\mathbf{X}_f \setminus \{X_{fk}\}]$.

The Shapley value of variable X_{fk} for individual s in family f is then defined as the average of its marginal contributions, where the average is taken over all possible orderings

¹³These properties are additivity and monotonicity. Additivity ensures that for a given set of covariates \mathbf{X}_f , the sum of the Shapley values equals the model’s prediction $\hat{g}(\mathbf{X}_f)$. Monotonicity guarantees that if a variable’s contribution increases or stays the same, its Shapley value will not decrease, regardless of the other inputs.

¹⁴Their key insight is that trees are particularly suited because moving down a path in a tree amounts to adding variables one by one through their splits. This structure makes it possible to track each variable’s contribution to the prediction without re-estimating the model for all possible subsets.

of the covariates:

$$\phi_{sf}(X_{fk}) = \sum_{S \subseteq \mathbf{X}_f \setminus \{X_{fk}\}} w(S) h(X_{fk}, S), \quad (3)$$

where $w(S) = \frac{|S|!(K-|S|-1)!}{K!}$. For example, if parental income has a Shapley value of -2 for a given child, this means that including parental income lowers her expected income rank by two on average, where the average is taken over all possible subsets of included covariates.

For each child, the Shapley values of all variables sum to the gap between the child's expected income and the population mean:

$$g(\mathbf{X}_f) - E[Y_{sf}] = \sum_{k=1}^{91} \phi_{sf}(X_{fk}). \quad (4)$$

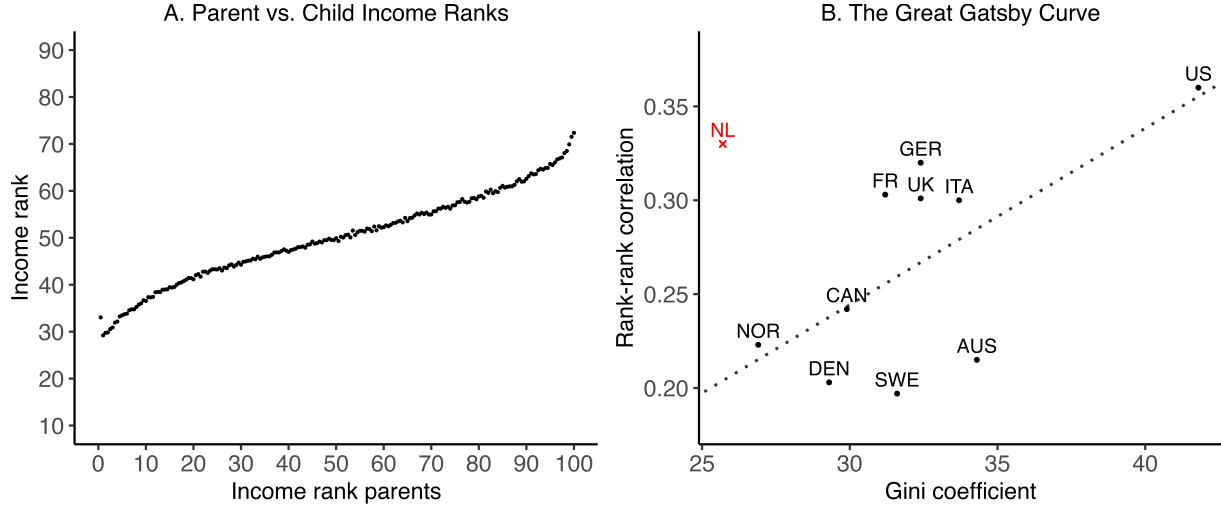
This additive decomposition of variables' contributions does not imply that the underlying conditional expectation function $g(\mathbf{X}_f)$ is assumed to be additive in these variables. When two variables interact, the marginal contribution of one depends on whether the other is included in the subset of variables considered. Because Shapley values average these marginal contributions over all possible subsets, they naturally account for interaction effects whenever the interacting variables appear together.

4 Main Results

4.1 Intergenerational Income Mobility in the Netherlands

I begin with a baseline analysis of intergenerational income mobility in the Netherlands. Figure 1 (a) presents a binscatter plot of children's income ranks relative to their parents' income ranks. The X-axis is divided into 200 bins, each containing roughly 8,500 children. The dots correspond to the mean household income rank of children given their parents' household income rank. Child income increases linearly between the 10th and the 90th

Figure 1: Intergenerational Income Mobility in the Netherlands



Notes: panel A presents a nonparametric scatter plot of mean income ranks versus parental income rank. The sample consists of all $N = 1,712,291$ Dutch children born between 1980 and 1989 (Table A1) for whom their own and parental income is not missing (96.2%). The sample exceeds the core sample in Figure 2 because it includes adopted children. The X-axis reports the parent income rank sorted into 200 equal-sized bins. The Y-axis reports the mean income rank within each bin. Panel B presents a cross-country comparison of Gini coefficients and rank-rank correlations. The dotted line shows the regression line obtained when regressing rank-rank correlations on Gini coefficients, excluding the Netherlands. The Gini coefficients are taken from the most recent estimates between 2018 and 2023 from the World Bank. The rank-rank correlations are computed by: Heidrich (2017) (Sweden), Helsø (2021) (Denmark), Deutscher and Mazumder (2020) (Australia), Bratberg et al. (2017) (Norway), Corak (2020) (Canada), Kenedi and Sirugue (2023) (France), Dodin et al. (2024) (Germany), Acciari et al. (2022) (Italy), Rohenkohl (2023) (the UK), and Davis and Mazumder (2024) (the US). To make estimates comparable, the Italian estimate corresponds to the one when adjusting for lifecycle bias, incomplete coverage of taxpayers, and tax evasion (reported on page 28). See Kenedi and Sirugue (2023) for a more detailed comparison of approaches.

income ranks but increases steeply at the tails of the parental income distribution.¹⁵ An OLS regression yields a slope coefficient of 0.33, indicating that a one-rank increase in parental income corresponds to a 0.33-rank increase in children’s income on average.¹⁶

While cross-country comparisons should be made with caution, Figure 1(b) suggests that the Netherlands ranks among OECD countries with strong persistence. Its rank-rank correlation is higher than in Sweden, Denmark, Australia, Norway, and Canada (0.20–0.24) and

¹⁵As noted before by Van Elk et al. (2024), there is some measurement error at the very bottom of the parental income distribution. This is because some wealthy parents report low income as a result of capital losses. Removing the bottom 0.5 percent of the sample does not affect the estimates much.

¹⁶This estimate exceeds recent estimates for the Netherlands from Van Elk et al. (2024), Manduca et al. (2024), and Boustan et al. (2025), who report estimates between 0.16 and 0.23. In Appendix B, I replicate their approaches and illustrate why my estimates differ.

more similar to France, Germany, Italy, the UK, and the US (0.30-0.36). This is somewhat surprising because the Netherlands has one of the lowest levels of income inequality, and lower inequality is often associated with lower persistence. Indeed, Panel B suggests that the Netherlands is an outlier on the ‘Great Gatsby Curve’ (Corak (2013)).

Appendix B reports additional, commonly used intergenerational mobility estimates to facilitate cross-country comparisons. Moreover, I vary the number of years over which parental income is measured and the timing of income measurement in parents’ and children’s lives. These robustness checks suggest that the estimate is robust to transitory measurement error and lifecycle bias.

4.2 Measuring Income Gaps with all Family Information

I next quantify income gaps across family backgrounds using all 91 family characteristics, and compare them to gaps based on parental income alone. For consistency, I repeat the analysis using parental income alone, this time using the same training and test data as the full model. The *income-only* model predicts a child’s income rank nonparametrically as the mean rank of children in the training data with the same parental income rank and year of birth. The *comprehensive* model is a tuned gradient-boosted decision tree using all variables from Table 1 and year of birth.

Figure 2 presents binscatter plots corresponding to each model, constructed as follows. Consider first the blue diamonds, which use predictions from the comprehensive model. I sort the test sample into 200 equal-sized bins, each containing approximately 1,700 children, based on their predicted income rank within their cohort. The Y-axis reports the mean observed income rank within each bin. Moving from left to right shows how mean income rises as family background becomes more favorable. The grey dots are constructed in the same way, except that children are sorted by predictions from the income-only model.¹⁷

Mean income rises more steeply under the comprehensive model. The comprehensive

¹⁷Because expected income is a monotonic function of parental income up to minor deviations (Figure 1), this is almost identical to sorting children by parental income rank.

model thus identifies favorable and unfavorable family backgrounds more accurately: at low predicted-income bins, children have lower observed income than under the income-only model, and at high predicted-income bins they have higher observed income. The higher accuracy is also reflected in the R^2 , which increases by 55 percent relative to the income-only model. The income-only model’s flatter curve and lower R^2 reflect a more limited ability to identify children with different expected income. Some low-income families are advantaged on other dimensions, and some high-income families are not. Unlike the comprehensive model, the income-only model is unable to distinguish these families.

The comprehensive model reveals strong and previously unseen income disparities at the tails. Among the 0.5 percent of children with the lowest expected income, the mean observed income rank is 17. Among the next 0.5 percent, it is 19.5—a difference of 2.5 ranks across less than one percent of the distribution.¹⁸ This pattern is similar at the very top: the 0.5 percent with the highest expected income reach a mean rank of 78, compared to 74.6 for the next 0.5 percent. Taken together, the overall income gap between the children from the least and most advantaged families is 61 ranks. This is equivalent to 146,000 euros in gross household income annually. The income-only model strongly underestimates this gap: under the income-only model, it equals 42 ranks (from 31 to 73), or 54,000 euros annually.¹⁹

Conditional income probabilities. Figure 2 reports only mean income ranks, yet there is substantial variation around this. To capture this, Figure 3 (a) presents a 200×5 matrix with the shares of children in each adult income quintile given expected income. In the absence of any intergenerational transmission, the share of children in each income quintile should be 20 percent. Instead, there are large differences. Among the 0.5 percent children with the lowest expected income, less than 1 percent are in the top income quintile, whereas 72 percent are in the bottom income quintile. Among the 0.5 percent children with the

¹⁸To put this in perspective, across the entire population, a half percentile ‘more advantaged’ family is associated with only a 0.21 ranks higher mean income.

¹⁹The annual gross household incomes corresponding to these ranks are: rank 17 = €48k, rank 78 = €194k, rank 31 = €74k, and rank 73 = €128k.

highest expected income, these shares are 64 and 5 percent, respectively.

To place the conditional probabilities in international perspective, panel B focuses on the 20 percent of children with the lowest expected income. Because similar measures using multiple family characteristics are not available in other countries, I compare their probabilities of reaching the top quintile (‘moving up’) or being in the bottom quintile (‘staying down’) to estimates based on parental income alone, both for the Netherlands and for other countries. Even using only parental income, the Netherlands already shows strong persistence at the bottom: the twenty percent least advantaged children are four times more likely to be in the bottom quintile than the top (36 vs. 9 percent).²⁰ Incorporating the full set of family background variables increases this ratio to seven (42 vs. 6 percent). These results indicate that for a substantial share of Dutch children, the chances of ‘moving up’ are low, and much lower than implied by parental income alone.

Robustness. Table A3 shows that explanatory power declines with smaller samples but stabilizes once at least 40 percent of the data are used. The same holds for the observed income gap between the 0.5 percent children with the lowest and highest predicted income. This suggests that downward bias due to insufficient sample size for training the machine learning model is unlikely.

Table A4 varies the number of years and ages at which child income is measured. The observed income gap between the 0.5 percent children with the lowest and highest predicted income remains stable across all specifications. Explanatory power attenuates when fewer years of income are used, but stabilizes once about five years of income are used.²¹ It also decreases somewhat when income is measured exclusively in the early 30s, but when I re-estimate the model using only incomes beyond age 32, then the overall estimate is very

²⁰In Figure B1, I provide the full 5×5 transition matrix based on parental income only.

²¹In intergenerational mobility regressions, classical measurement error in child income does not bias the coefficient estimate. It only inflates the standard error. However, when estimating explanatory power, such left-hand side measurement error does matter. Reducing measurement error lowers the variance of child income, which in turn affects the explanatory power of the regression.

similar to the main results specification. This indicates that the influence of attenuation or life-cycle bias is likely minimal.

Finally, I assess the importance of the missing education records. I first train the model on the subset of children whose parents' education is observed ($n = 1,105,905$, $R^2 = 17.7$). I then re-train the model on the same sample after removing all education variables for both parents and extended family ($R^2 = 17.6$). The resulting drop in R^2 is only 0.1 percentage point, indicating that the remaining variables already capture most of the educational variation across families. This suggests that the explanatory power of the model would increase only marginally if complete education data were available.

4.3 What Characterizes Family (Dis)Advantage?

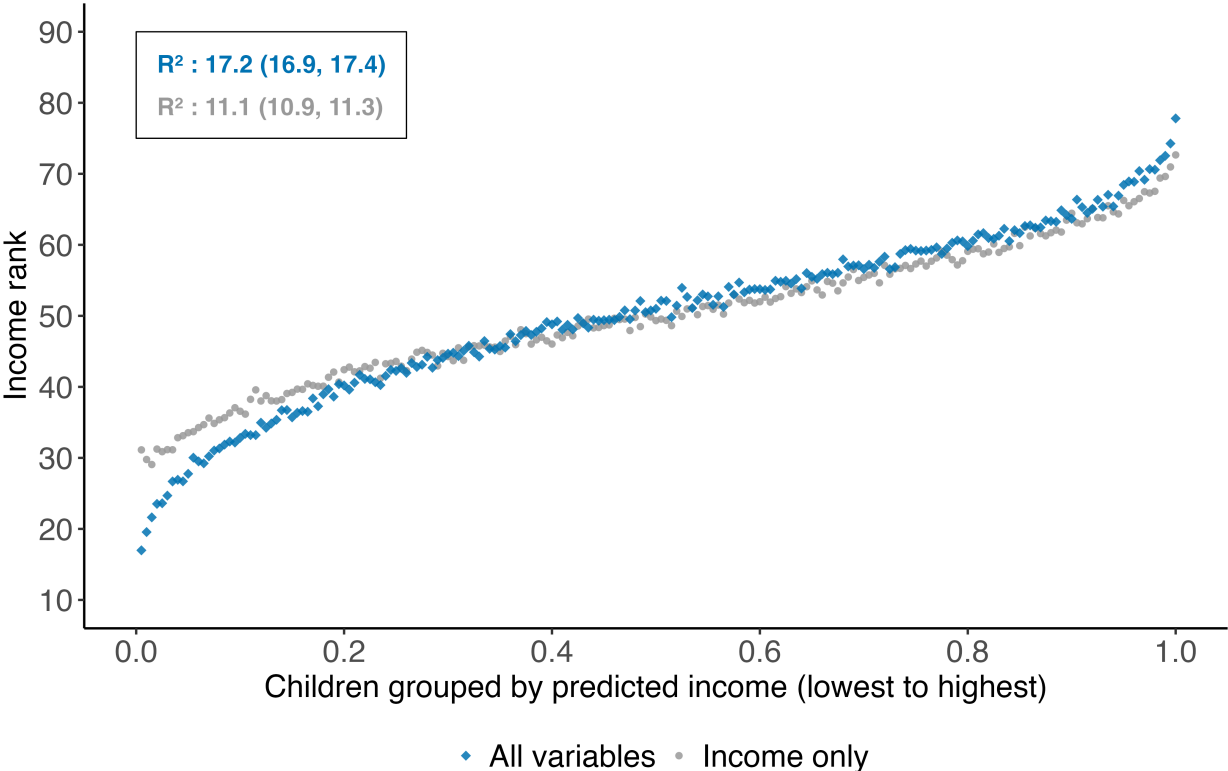
Distribution of Shapley values. Which family background characteristics drive the expected income differences between children? Answering this requires attributing changes in expected income to specific variables. I do this using the Shapley decomposition, which assigns to each variable its average marginal contribution to a child's expected income rank across all subsets of the remaining 90 variables.²² This approach produces 91 Shapley values per child that sum exactly to the gap between the child's expected income rank and the population mean of 50 (Equation 4).

To summarize the results compactly, I sum the Shapley values of variables belonging to the same category in Table 1. This yields nine indices per child, each of which measures the total contribution of all variables in its category. A parental income index of -5, for example, means that the parental income variables jointly contribute -5 ranks to the gap between the child's expected income rank and the population mean of 50.

Figure 4 reports the distribution of each index across all children in the test data. The distribution of Shapley values for the parental income variables is wider than for any other

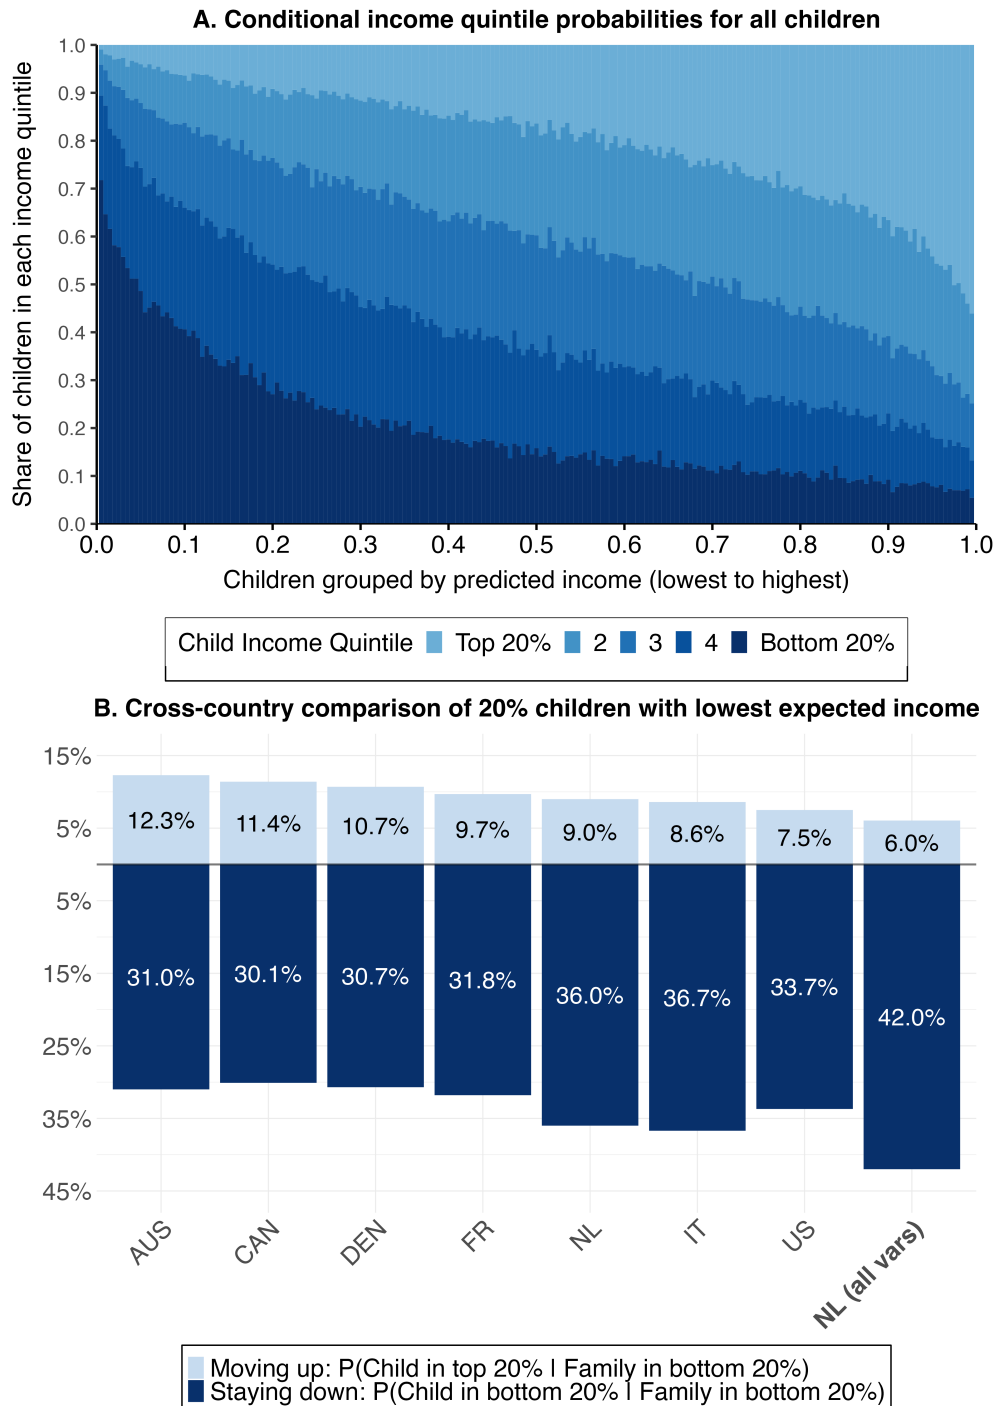
²²For example, a Shapley value of -3 for parental income rank means that adding parental income rank to the conditioning variables lowers the child's expected income by 3 ranks, where the average is taken over all possible subsets of other conditioning variables. See Section 3 for a detailed explanation.

Figure 2: Predicting Child Income with Detailed Parental Information



Notes: this Figure presents binscatter plots of income ranks for 341,107 children in the test data, who are sorted into bins based on their predicted income rank according to two models. Both models are trained to predict children’s income ranks using the same training sample of 1,364,425 children but include different explanatory variables. The blue graph is constructed as follows: (i) predict the income ranks of all children in the test data using the model with all explanatory variables, (ii) rank the predictions from low (0) to high (1) within a child’s cohort, (iii) sort all children into 200 equal-sized bins based on their ranking, and (iv) calculate the average income ranks within each bin. The grey graph is constructed similarly using the predictions from the model that uses parents’ income only. Confidence intervals for the R^2 are bootstrapped from the test data using 599 draws.

Figure 3: Child Income Distributions Across Family Backgrounds



Notes: panel A shows the share of children in each income quintile. The sample contains all 341,107 children from the test data, and they are grouped into 200 equally sized bins according to their predicted income rank from the comprehensive model (as in Figure 2). Panel B shows the share of children in the top (transparent) or bottom (dark) income quintile. The first seven bars represent children from families in the bottom 20 percent of the parental income distribution. Estimates for other countries are taken from the studies listed below Figure 1. The final bar represents the 20 percent of children with the lowest predicted incomes from the comprehensive model (the bottom twenty percent in panel A).

category. This means that variables from this category make large contributions to the expected income of many children, shifting it both upward (for children with favorable parental income variables) and downward (for children with unfavorable parental income variables). The histograms for parental wealth and extended-family outcomes are also relatively wide. The other histograms are more centered around zero, indicating small contributions to expected income from these variables for most children.

The figures also report the share of the total variation in expected income that each of these indices explains, indicated by S .²³ Consistent with their wide distributions, parental income, parental wealth, and extended-family outcomes jointly account for 78 percent of the variation in expected income.²⁴ This underscores the importance of accurately measuring economic resources and extended family links for quantifying family-driven income inequality.

Shapley values for children at the tails. The dashed lines in Figure 4 report the mean of each index for the 0.5 percent of children with the lowest and highest expected incomes. They reveal an interesting asymmetry: the high expected income of the most advantaged children mostly reflects the contributions of the same three factors that account for most of the cross-population variation, whereas the low expected income of the most disadvantaged children reflects negative contributions from all family background dimensions.


The children at the top have a mean expected income rank of 78, so their Shapley values sum to +28. The dashed lines show that parental income (+14), parental wealth (+5), and extended-family outcomes (+5) account for 24 of these ranks; the remaining variables jointly contribute only 4 ranks.

The children at the bottom have a mean expected rank of 16, so their Shapley values sum to -34. Parental income (-12), parental wealth (-6), and extended-family outcomes

²³The share S is computed using the decomposition rule of Shorrocks (1982), which applies to outcomes that are an additive sum of factors. For each Shapley index ϕ_k , $S(\phi_k) = \text{Cov}(\phi_{ik}, \hat{Y}_i) / \text{Var}(\hat{Y}_i)$, where ϕ_{ik} equals the Shapley index for child i corresponding to category $k \in \{1, \dots, 9\}$. The decomposition applies here since $\hat{Y}_i - E[Y] = \sum_{k=1}^9 \phi_{ik}$ is an additive sum of contributions.

²⁴Figure A4 plots the Shapley values of the 30 most predictive variables individually. Eight of the top ten relate to parental or extended-family income and wealth; the remaining extended-family outcomes (health, crime, and education) contribute little.

Figure 4: Measuring the Contribution of Family Background Variables to Expected Income with Shapley Values



Notes: each panel shows the distribution of a Shapley index across the 341,107 test-data children. Each index sums the Shapley values of variables in one family-background category (Table 1), computed using Lundberg et al. (2020). S is the share of variation in predicted income explained by the category (Shorrocks (1982)). Dashed lines mark the mean index for the 0.5% of children with the lowest and highest predicted income.

(-6) contribute similarly in magnitude to what they contribute at the top, accounting for -24 ranks. The other dimensions, however, contribute much more than at the top: family structure (-5), education (-2), crime (-2), occupation (-1), health expenditures (-1), and migration background (-1) together add another -12 ranks. These variables thus account for a non-trivial share of the very low expected income at the bottom, even though their contributions to the cross-population variation are relatively small.

Table A2 reports descriptive statistics of family characteristics across the expected income distribution. The raw data confirm the patterns from the Shapley decomposition: the children at the bottom have parents with low income and wealth and who are often young, separated, and minimally educated; their parents are frequently suspected of crimes and have high health expenditures; and their aunts and uncles also have low income and wealth. Children at the top have favorable characteristics on all of these dimensions.

These relationships are associations and do not necessarily reflect causal effects. Nevertheless, understanding what predicts low income can inform policy. The results show that the identification of the most disadvantaged children relies on the cumulative contribution of multiple family-background variables. This implies that support programs selecting families on a single characteristic, such as income, education, or migration background, face two problems. They miss children whose families are disadvantaged on many dimensions but not on the targeting variable, and they include children who are not disadvantaged on other dimensions. Combining several characteristics can substantially improve this targeting.

4.4 Additional Results

Sibling correlation. I estimate a sibling correlation in income of 0.301 (Table B1). As discussed in Section 3, this provides an upper bound on the explanatory power of any predictive model that solely includes variables that are equal between siblings, as in this paper. The explanatory power of the comprehensive model is over half of this correlation (0.172/0.301). The remainder of the sibling correlation may be explained by other shared factors, such

as unobserved parental characteristics, community influences, shocks, or spillovers, that are uncorrelated with the included variables.

Gender differences. Figure A1 presents results from predictive models trained to predict sons' and daughters' household income ranks separately. I also present results using personal income ranks in Figure A2 to abstract away from household formation considerations. The explanatory power for predicting household income ranks is similar between genders, and for predicting personal income ranks, it is somewhat higher for daughters.

Predicting education and violent crime. Figure A3 presents results for children's education and violent crime. These outcomes are interesting in their own right, but they also offer a way to assess whether richer family information adds more or less value when predicting outcomes other than income ranks. The explanatory power for education rises from 12.6 percent to 26.1 percent, an increase of 107 percent. For violent crime it rises from 3.9 percent to 10.5 percent, an increase of 169 percent. Both gains far exceed the 55 percent increase for income ranks. This shows that broader family information can be even more relevant for quantifying disparities in other outcomes.

Functional form. A straightforward OLS model, which includes all variables linearly, achieves an explanatory power of 15.7 percent.²⁵ The mean income gap between the children with the 0.5 percent lowest and highest predicted income ranges from ranks 19 to 75. This is quite close to the results of the comprehensive model, suggesting that incorporating a broader range of family characteristics is more critical than allowing for complex interactions and non-linearities. A downside of OLS regression is that it does not support the computation of individual-specific Shapley values.²⁶

²⁵Coefficient estimates are available upon request.

²⁶While in theory this is possible by estimating new OLS regressions for each possible subset of explanatory variables, in practice this is infeasible with the large number of explanatory variables.

Analysis with adoptees. Appendix C applies the income predictions to a sample of international adoptees to quantify the role of post-birth factors in driving the observed income gaps. The results indicate that being raised from infancy in a more advantaged family increases adoptees' income, but by 70 percent less than what is expected based on own-birth children's outcomes. A Shapley value analysis shows that this attenuation is stronger for parental income and extended family outcomes than for wealth and family structure, suggesting that income gaps along the latter dimensions reflect relatively stronger differences in post-birth factors.

5 Measuring Income Gaps Across Neighborhoods

A growing literature documents large geographic differences in the expected income of children from *low-income* families (Heidrich (2017), Deutscher and Mazumder (2020), Corak (2020), Alesina et al. (2021), Acciari et al. (2022), Kenedi and Sirugue (2023), Chetty et al. (2026)). This section shows that such geographic variation reduces substantially when measuring family disadvantage using all family information.

Child addresses are first observed in 1995, when the earliest cohort is fifteen years old. I am able to link 98 percent of the analysis sample to the 2,829 neighborhoods in which they were registered at age fifteen. Neighborhoods follow the classification of Statistics Netherlands and are comparable to United States census tracts. Their average population is about 4,900 individuals. For reference, among children born in 1995 with a complete residential history up to age eighteen, the median child spent seventeen of those years in the neighborhood where they were registered at age fifteen, and on average children spend 14.3 years (80%) of these years in this neighborhood. This shows that the neighborhood at age fifteen provides a good measure of where children spent most of their childhood.

Upward mobility measures. I follow Chetty et al. (2026) to estimate differences in inter-generational mobility across neighborhoods. They define *Absolute Upward Mobility* (AUM) for neighborhood n as the expected income rank of a child in the 25th percentile of the parental income distribution: $\text{AUM}(n) = E[Y_{sf}|Y_f = 25, n]$. Because there are often too few children in a neighborhood to estimate this nonparametrically, they use a univariate regression whose functional form is chosen based on estimates at the national level:

$$Y_{sf} = \alpha_n + \beta_n \hat{Y}(Y_f) + e_{sf}, \quad (5)$$

where $\hat{Y}(Y_f)$ nonparametrically estimates $E[Y_{sf}|Y_f]$ based on the national-level relationship between Y_{sf} and Y_f .²⁷ This specification summarizes the conditional expectation function $E[Y_{sf}|Y_f, n]$ in each neighborhood using just two parameters, α_n and β_n .

I generalize this approach to include additional family information. Let \hat{Y}_{25} be the 25th percentile of the national expected income distribution based on all observable family characteristics: $\hat{Y}_{25} = \{Y : P(\hat{Y}(\mathbf{X}_f) \leq Y) = 0.25\}$. I define *Multidimensional Absolute Upward Mobility* (MAUM) as the expected income of children who grew up in neighborhood n and who are in the 25th percentile of the national expected income distribution: $\text{MAUM}(n) = E[Y_{sf}|\hat{Y}(\mathbf{X}_f) = \hat{Y}_{25}, n]$. This measure equals AUM when parental income is the only observed family background characteristic, but they differ when there are additional family background variables that predict child income.²⁸ To estimate MAUM, I use the same univariate regression as in Equation 5, but replace the predictions $\hat{Y}(Y_f)$ by predictions $\hat{Y}(\mathbf{X}_f)$:

$$Y_{sf} = \delta_n + \gamma_n \hat{Y}(\mathbf{X}_f) + \nu_{sf}, \quad (6)$$

²⁷Chetty et al. (2026) also condition on child gender and race since they study these groups separately. I omit these variables because I pool all observations.

²⁸This follows because children whose parents rank 25th in the national income distribution rank 25th themselves in the expected income distribution if parental income is the only observed characteristic.

where $\hat{Y}(\mathbf{X}_f)$ estimates the national-level conditional expectation function $E[Y_{sf}|\mathbf{X}_f]$.

I use the same predictive models as in Figure 2 to generate predictions $\hat{Y}(Y_f)$ and $\hat{Y}(\mathbf{X}_f)$. I include all children in the core sample, including those originally in the training sample for the predictive model, to increase the number of observations in each neighborhood. A cross validation procedure ensures that predictions for these children remain out of sample.²⁹ I then estimate $AUM(n)$ and $MAUM(n)$ by running regressions 5 and 6 for each neighborhood and plugging in $\hat{\alpha}_n$, $\hat{\beta}_n$, $\hat{\delta}_n$, $\hat{\gamma}_n$, $Y_f = 25$, and $\hat{Y}(\mathbf{X}_f) = \hat{Y}_{25}$. When computing the standard deviations of the neighborhood-level mobility estimates, I weight neighborhoods by the number of children with below-median parental income.

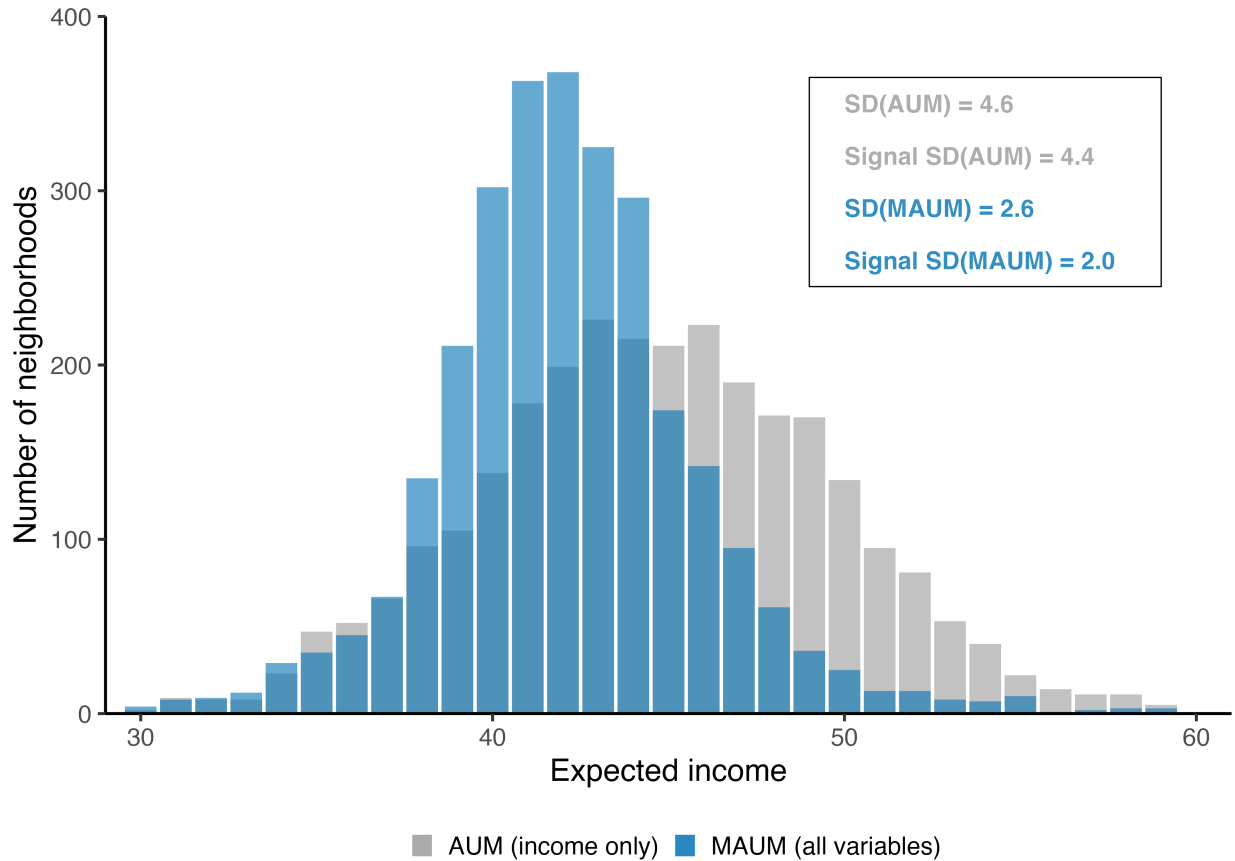
Results. Figure 5 shows substantial variation in absolute upward mobility, with estimates ranging from 30 to 60 ranks. The standard deviation is 4.6 ranks. Taken at face value, this suggests that a one standard deviation higher upward mobility neighborhood increases expected income by 4.6 ranks. However, part of the variation in the estimates reflects noise induced by sampling uncertainty, since there are on average 590 observations per neighborhood. I use the estimates' standard errors to conduct a standard signal-noise decomposition, and estimate that the signal standard deviation equals 4.4 ranks.³⁰ This is somewhat lower than the variation across US census tracts, where the signal standard deviation in absolute upward mobility equals 6.2 ranks (Chetty et al. (2026)).

A sizable share of the variation in absolute upward mobility reflects sorting. To show this, I select all children with low parental income (between the 20th and 30th percentile) and regress their predicted income $\hat{Y}(\mathbf{X}_f)$ on neighborhood absolute upward mobility $A\hat{U}M$, controlling for parental income. Absent sorting, the coefficient should be zero: predicted income depends only on family characteristics, so it should be unrelated to neighborhood mobility. Instead, Table A7 shows a coefficient of 0.67. This implies that based on children's

²⁹Table A5 shows that the R^2 of the predictions for the full sample is 16.9 percent. This is in line with the results in Figure 2. See Table A5 for a step-by-step explanation of the cross-validation procedure.

³⁰Specifically, I estimate the signal variance by subtracting the mean squared standard error from the sample variance of the estimates, and then take the square root to obtain the signal standard deviation.

Figure 5: Variation in Absolute Upward Mobility Across Neighborhoods



Notes: this Figure presents histograms of neighborhood-specific upward mobility estimates for 2,829 neighborhoods. The Absolute Upward Mobility (AUM) estimates are obtained by the fitted values of regression 5 after plugging in $Y_f = 25$. The Multidimensional Absolute Upward Mobility (MAUM) estimates are obtained by the fitted values of regression 6 after plugging in $\hat{Y}(\mathbf{X}_f) = \hat{Y}_{25}$. Estimates below 30 and above 60 (1 percent of estimates) are dropped. Standard deviation estimates are weighted by the number of children with below-median parental income in each neighborhood. Signal standard deviation estimates are computed by subtracting the weighted average squared standard error from the weighted sample variance of estimated upward mobility, then taking the square root.

family characteristics beyond parental income we would already expect children in one-rank higher upward mobility neighborhoods to have 0.67 ranks higher income.

The multidimensional measure already accounts for such expected income differences. As a result, the multidimensional estimates show much lower dispersion.³¹ Their standard deviation is 2.6 ranks, which is 43 percent lower than the estimates that use parental income only. This drop does not reflect lower sampling uncertainty, since the standard errors are very similar. The lower dispersion must therefore arise entirely from a decline in the underlying signal, which must then exceed 43 percent. The signal noise decomposition confirms this: the signal standard deviation is 2 ranks, which is 55 percent lower than the signal in the estimates using parental income only. This shows that variation in expected income of disadvantaged children across neighborhoods is far smaller when family disadvantage is proxied by a low expected income based on all family information rather than parental income alone.

Policymakers and families care not only about levels of upward mobility but also about the ranking of neighborhoods. Figure A5 shows that the upward mobility ranking of all neighborhoods can change substantially with the multidimensional measure. Neighborhoods at the same percentile under the income-only measure are, on average, 20 percentiles apart under the multidimensional measure.³² To illustrate, take the 142 neighborhoods between the 10th and 15th percentile under the income only measure. In the new ranking, the lowest 10 are among the 7 percent lowest mobility neighborhoods, whereas the highest 10 rise above the 57th percentile.

Overall, the reduced standard deviation of 2 ranks leaves a relatively modest role for neighborhoods in driving income gaps. It implies that, holding observed family background constant, a one standard deviation higher mobility neighborhood raises expected income by 2 ranks. By comparison, holding the neighborhood constant, a one standard deviation more advantaged family background raises expected income by 11.3 ranks.³³ This shows that in-

³¹The mean is also lower because $E[Y_{sf}|\hat{Y}(\mathbf{X}_f) = \hat{Y}_{25}] < E[Y_{sf}|Y_f = 25]$ (Figure 2).

³²This is measured by the within-group standard deviation of the new ranks, where groups are neighborhoods that fall in the same percentile based on the old ranking.

³³The standard deviation of $\hat{Y}(\mathbf{X}_f)$ is 11.8 ranks. Within neighborhoods, a one rank increase in $\hat{Y}(\mathbf{X}_f)$

come gaps are considerably larger across family backgrounds than across neighborhoods.

Robustness. Table A6 reports several alternative estimates. The reduction in dispersion is similar when I focus on children at the 5th or 75th percentile of the predicted income distribution. Table A6 also shows the variation in $\hat{\beta}_n$ and $\hat{\gamma}_n$, which capture *relative* inter-generational mobility. The dispersion in these estimates falls by 52 percent once I use all family information. I then repeat the analysis at the municipality level, which provides less granularity but is also less affected by movers.³⁴ Also at this higher level of aggregation, dispersion falls by 53 percent. These results show that the decline in dispersion is robust across multiple mobility measures and a higher level of aggregation.

Next, I assess two types of functional form restrictions. First, I estimate regressions 5 and 6 with second- or third-order polynomials to allow for nonlinearities in the relationship between predicted and observed child income. Table A6 shows that signal dispersion is similar under these specifications, suggesting that nonlinearities in the conditional expectation functions do not meaningfully change conclusions.

A second concern is the single-index restriction. MAUM compares children across neighborhoods who share the same value of $\hat{Y}(\mathbf{X}_f)$, but many combinations of \mathbf{X}_f can produce the same predicted income. If the relative importance of different elements of \mathbf{X}_f varies across neighborhoods, cross-neighborhood variation in $E[Y_{sf}|\hat{Y}(\mathbf{X}_f), n]$ may differ from variation in the ideal measure $E[Y_{sf}|\mathbf{X}_f, n]$. Estimating the latter within each neighborhood is not feasible given the number of predictors and the sample sizes involved.

Instead, I test whether the results are robust to a more flexible specification by replacing predicted income with the nine Shapley-value indices from Section 4. Each index $\phi_k(\mathbf{X}_f)$ represents a different family background dimension and measures its contribution to a child's predicted income. For example, a family with a positive income index but a negative wealth

increases income by 0.96 ranks (Table A5). Keeping the neighborhood constant, a one standard deviation increase in $\hat{Y}(\mathbf{X}_f)$ therefore raises expected income by $0.96 \times 11.8 = 11.3$ ranks.

³⁴The mean municipality population is 44,000 persons. Children spend on average 87 percent of their first eighteen years in the municipality where they were registered at age fifteen.

index has favorable income characteristics (that contribute positively to the child’s expected income) but unfavorable wealth characteristics (that contribute negatively to the child’s expected income). Because the indices sum to the total prediction, they are jointly as predictive of child income at the national level as the full set of family background variables.³⁵ However, dimensionality reduces from 91 variables to nine. This makes it feasible to estimate a separate slope for each dimension within each neighborhood.

I then estimate the following regression in each neighborhood:

$$Y_{sf} = \kappa_n + \sum_{k=1}^9 \gamma_{n,k} \phi_k(\mathbf{X}_f) + \eta_{sf}. \quad (7)$$

Equation 6 nests this regression via the restriction $\gamma_{n,1} = \gamma_{n,2} = \dots = \gamma_{n,9}$.³⁶ Equation 7 relaxes this restriction by allowing each dimension to have its own neighborhood-specific slope. I conduct an F-test in each neighborhood for the null hypothesis that the restriction holds. I then compute an alternative measure of multidimensional upward mobility by plugging $\hat{\kappa}_n, \hat{\gamma}_{n,1}, \dots, \hat{\gamma}_{n,9}$ into equation 7 and fixing each index to its average value among children at the 25th percentile of $\hat{Y}(\mathbf{X}_f)$.³⁷

The restriction is rejected at the five percent level in ten percent of neighborhoods, only five percentage points above the rate expected under the null (Table A8). This low rejection rate indicates that, in the large majority of neighborhoods, relaxing the single-index restriction is either unnecessary (because the restriction holds) or unhelpful (because the sample is too small to estimate the additional parameters with sufficient precision). Most importantly, the signal SD of the unrestricted measure is 2.2 ranks, very close to MAUM’s 2.0 ranks (Figure 5), while the raw standard deviation rises from 2.6 to 3.6. Allowing for additional flexibility thus leaves the underlying cross-neighborhood dispersion in upward mobility unchanged while substantially reducing the precision of neighborhood-specific estimates.³⁸

³⁵See Table C2 column 1.

³⁶This follows by plugging in $\hat{Y}(\mathbf{X}_f) = E[Y] + \sum_{k=1}^9 \phi_k(\mathbf{X}_f)$ into equation 6.

³⁷This ensures that the measure is evaluated at the same predicted income level as MAUM.

³⁸Table A8 also reports municipality level results, where power is much larger. The F-test rejects in 19

Discussion. Prior studies often use (quasi)random moves to estimate the share of variation in upward mobility that reflects family sorting (Chetty et al. (2016), Chetty and Hendren (2018), Kawano et al. (2024)). Chetty et al. (2026), for example, attribute about 40 percent of the variation across US census tracts to sorting. The appeal of using (quasi)random moves is that it nets out all family background differences, observed and unobserved. A limitation is that the number of movers per neighborhood is typically too small to yield reliable new neighborhood-specific estimates, so both the upward mobility level and the ranking of specific neighborhoods remain uncertain. The results above show that adjusting for a broad set of *observable* family background characteristics can already correct for large sorting patterns while still using the full sample. The new estimates should be closer to the ideal measures that policymakers need to identify neighborhoods that offer weak income prospects for otherwise comparable children.³⁹

Lastly, Cholli et al. (2024) apply a related approach to measure upward mobility in Danish data. They generalize a standard intergenerational mobility regression by including observable family and social characteristics and they use a control function approach to additionally address selection on unobservables. Both their approach and mine explore the role of family sorting in shaping cross-neighborhood differences in upward mobility and construct new estimates that adjust for it.⁴⁰ However, the two approaches differ in how they incorporate family information and in what their neighborhood-specific mobility parameters measure. Appendix D discusses these conceptual differences in more detail.

percent of municipalities, and the signal standard deviation is identical to MAUM's (1.6 ranks). This shows that even with larger samples, the single-index restriction is reasonable in most municipalities, and relaxing it does not change the conclusions.

³⁹Future work could combine the multidimensional approach with (quasi)random moves to assess how much of the remaining variation still reflects unobserved family background differences. This is not possible here because I do not observe addresses early enough.

⁴⁰They report reductions in the signal standard deviation of absolute upward mobility of 25 percent at the municipality level and 29 percent at the parish level after adding family background controls. See Appendix D for a more detailed comparison of results.

6 Conclusion

This paper provides new insights into the income prospects of the least and most advantaged children by combining a multidimensional approach with Dutch administrative data containing the most comprehensive family information studied to date. The main contribution of this paper is to offer an analysis that is simultaneously *broad*—covering many family background characteristics—and *local*—presenting granular distributional insights into both the size of income gaps and the decomposition of their sources across the full population.

This approach reveals that the mean income gap between the least and most advantaged children spans 61 ranks, from rank 17 to 78, and that a conventional income-based measure understates this gap by 19 ranks. A novel Shapley value decomposition shows that parental and extended family income and wealth account for the largest share of the income disparities, but that accurate identification of the most disadvantaged children requires more information. Their low expected income reflects the cumulative negative contribution of many family characteristics simultaneously.

I next use the multidimensional income predictions to produce new estimates of neighborhood upward mobility that adjust for differences in many observable family characteristics rather than parental income alone. Variation in upward mobility falls by over half under this measure, and neighborhood mobility rankings change substantially. Policymakers can use the new estimates to identify neighborhoods where children with similarly disadvantaged observed family backgrounds face persistently lower income prospects.

Advances in data quality, computing power, and statistical methods continue to expand the scope for empirical work on intergenerational mobility. This paper illustrates how many family characteristics can be analyzed jointly with modern machine learning methods and demonstrates the insights that this can yield for the measurement of economic inequality across family backgrounds and neighborhoods. As more comprehensive data become available, future research can use this approach to deepen our understanding of the complex process through which economic status is transmitted across generations.

References

- Acciari, Paolo, Alberto Polo, and Giovanni L. Violante.** 2022. “And Yet It Moves: Intergenerational Mobility in Italy.” *American Economic Journal: Applied Economics* 14 (3): 118–163.
- Adermon, Adrian, Gunnar Brandén, and Martin Nybom.** 2025. “The Relationship between Intergenerational Mobility and Equality of Opportunity.” IFAU Working Paper No. 2025:2.
- Adermon, Adrian, Mikael Lindahl, and Mårten Palme.** 2021. “Dynastic Human Capital, Inequality, and Intergenerational Mobility.” *American Economic Review* 111 (5): 1523–1548.
- Alesina, Alberto, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou.** 2021. “Intergenerational mobility in Africa.” *Econometrica* 89 (1): 1–35.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso.** 2018. “Intergenerational Mobility and Preferences for Redistribution.” *American Economic Review* 108 (2): 521–554.
- Althoff, Lukas, Harriet Brookes Gray, and Hugo Reichardt.** 2025. “America’s Rise in Human Capital Mobility.” CESifo Working Paper 11729.
- Bingley, Paul, and Lorenzo Cappellari.** 2019. “Correlation of Brothers’ Earnings and Intergenerational Transmission.” *The Review of Economics and Statistics* 101 (2): 370–383.
- Blundell, Jack, and Erling Risa.** 2019. “Income and Family Background: Are We Using the Right Models?” SSRN Working Paper.
- Boustan, Leah, Mathias Fjællegaard Jensen, Ran Abramitzky et al.** 2025. “Intergenerational Mobility of Immigrants in 15 Destination Countries.” NBER Working Paper No. 33558.
- 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.” *The Scandinavian Journal of Economics* 119 (1): 72–101.
- Brunori, Paolo, Francisco H.G. Ferreira, Guido Neidhöfer, and UNU-WIDER.** 2023. “Inequality of Opportunity and Intergenerational Persistence in Latin America.” WIDER Working Paper 2023.
- Brunori, Paolo, Francisco H.G. Ferreira, and Pedro Salas-Rojo.** 2024. “Inherited Inequality: A General Framework and a ‘Beyond-Averages’ Application to South Africa.” IZA Discussion Paper No. 17203.
- Chadwick, Laura, and Gary Solon.** 2002. “Intergenerational Income Mobility Among Daughters.” *American Economic Review* 92 (1): 335–344.
- Chang, Yoosoon, Steven N. Durlauf, Bo Hu, and Joon Park.** 2025. “Accounting for Individual-Specific Heterogeneity in Intergenerational Income Mobility.” NBER Working Paper 33349.
- Chetty, Raj, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter.** 2026. “The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility.” *American Economic Review* 116 (1): 1–51.
- Chetty, Raj, and Nathaniel Hendren.** 2018. “The Impacts of Neighborhoods on In-

- tergenerational Mobility I: Childhood Exposure Effects.” *The Quarterly Journal of Economics* 133 (3): 1107–1162.
- Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz.** 2016. “The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment.” *American Economic Review* 106 (4): 855–902.
- 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.” *The Quarterly Journal of Economics* 129 (4): 1553–1623.
- Cholli, Neil A., Steven N. Durlauf, Rasmus Landersø, and Salvador Navarro.** 2024. “Understanding the Heterogeneity of Intergenerational Mobility across Neighborhoods.” NBER Working Paper 33035.
- Collado, M Dolores, Ignacio Ortuño-Ortín, and Jan Stuhler.** 2023. “Estimating Intergenerational and Assortative Processes in Extended Family Data.” *The Review of Economic Studies* 90 (3): 1195–1227.
- Corak, Miles.** 2013. “Income Inequality, Equality of Opportunity, and Intergenerational Mobility.” *Journal of Economic Perspectives* 27 (3): 79–102.
- Corak, Miles.** 2020. “The Canadian Geography of Intergenerational Income Mobility.” *The Economic Journal* 130 (631): 2134–2174.
- Davis, Jonathan MV, and Bhashkar Mazumder.** 2024. “The Decline in Intergenerational Mobility after 1980.” *Review of Economics and Statistics* 1–47.
- Deutscher, Nathan, and Bhashkar Mazumder.** 2020. “Intergenerational Mobility across Australia and the Stability of Regional Estimates.” *Labour Economics* 66 101861.
- Deutscher, Nathan, and Bhashkar Mazumder.** 2023. “Measuring Intergenerational Income Mobility: A Synthesis of Approaches.” *Journal of Economic Literature* 61 (3): 988–1036.
- Dodin, Majed, Sebastian Findeisen, Lukas Henkel, Dominik Sachs, and Paul Schüle.** 2024. “Social Mobility in Germany.” *Journal of Public Economics* 232 105074.
- Eshaghnia, Sadegh, James J Heckman, and Rasmus Landersø.** 2025. “The impact of the level and timing of parental resources on child development and intergenerational mobility.” *Journal of Labor Economics* 43 (S1): S269–S301.
- Eshaghnia, Sadegh, James J. Heckman, Rasmus Landersø, and Rafeh Qureshi.** 2022. “Intergenerational Transmission of Family Influence.” NBER Working Paper 30412.
- Fagereng, Andreas, Magne Mogstad, and Marte Rønning.** 2021. “Why Do Wealthy Parents Have Wealthy Children?” *Journal of Political Economy* 129 (3): 703–756.
- Friedman, Jerome H.** 2001. “Greedy Function Approximation: A Gradient Boosting Machine.” *The Annals of Statistics* 29 (5): 1189–1232.
- Grinsztajn, Léo, Edouard Oyallon, and Gaël Varoquaux.** 2022. “Why Do Tree-Based Models Still Outperform Deep Learning on Tabular Data?” arXiv Working Paper 2207.08815.
- Haider, Steven, and Gary Solon.** 2006. “Life-Cycle Variation in the Association between Current and Lifetime Earnings.” *American Economic Review* 96 (4): 1308–1320.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman.** 2001. *The Elements of Statistical Learning*. Springer Series in Statistics, New York, NY, USA: Springer New York Inc..
- Heidrich, Stefanie.** 2017. “Intergenerational Mobility in Sweden: A Regional Perspective.”

- Journal of Population Economics* 30 (4): 1241–1280.
- Helsø, Anne-Line.** 2021. “Intergenerational Income Mobility in Denmark and the United States*.” *The Scandinavian Journal of Economics* 123 (2): 508–531.
- Holmlund, Helena, Mikael Lindahl, and Erik Plug.** 2011. “The Causal Effect of Parents’ Schooling on Children’s Schooling: A Comparison of Estimation Methods.” *Journal of Economic Literature* 49 (3): 615–651.
- Kawano, Laura, Bruce Sacerdote, William L Skimmyhorn, and Michael Stevens.** 2024. “On the Determinants of Young Adult Outcomes: Impacts of Randomly Assigned Neighborhoods for Children in Military Families.” NBER Working Paper 32674.
- Kenedi, Gustave, and Louis Sirugue.** 2023. “Intergenerational Income Mobility in France: A Comparative and Geographic Analysis.” *Journal of Public Economics* 226 104974.
- Lundberg, Scott M., Gabriel Erion, Hugh Chen et al.** 2020. “From Local Explanations to Global Understanding with Explainable AI for Trees.” *Nature machine intelligence* 2 (1): 56–67.
- Lundberg, Scott M, and Su-In Lee.** 2017. “A Unified Approach to Interpreting Model Predictions.” In *Advances in Neural Information Processing Systems*, Volume 30.
- Manduca, Robert, Maximilian Hell, Adrian Adermon et al.** 2024. “Measuring Absolute Income Mobility: Lessons from North America and Europe.” *American Economic Journal: Applied Economics* 16 (2): 1–30.
- Mazumder, Bhashkar.** 2005. “Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data.” *The Review of Economics and Statistics* 87 (2): 235–255.
- Nybom, Martin, and Jan Stuhler.** 2017. “Biases in Standard Measures of Intergenerational Income Dependence.” *The Journal of Human Resources* 52 (3): 800–825.
- Ramos, Xavier, and Dirk Van de Gaer.** 2016. “Approaches to Inequality of Opportunity: Principles, Measures and Evidence.” *Journal of Economic Surveys* 30 (5): 855–883.
- Roemer, John E., and Alain Trannoy.** 2016. “Equality of Opportunity: Theory and Measurement.” *Journal of Economic Literature* 54 (4): 1288–1332.
- Rohenkohl, Bertha.** 2023. “Intergenerational Income Mobility: New Evidence from the UK.” *The Journal of Economic Inequality* 21 (4): 789–814.
- Sacerdote, Bruce.** 2007. “How Large Are the Effects from Changes in Family Environment? A Study of Korean American Adoptees.” *The Quarterly Journal of Economics* 122 (1): 119–157.
- Sacerdote, Bruce.** 2011. “Chapter 1 - Nature and Nurture Effects On Children’s Outcomes: What Have We Learned From Studies of Twins And Adoptees?” In *Handbook of Social Economics*, edited by Benhabib, Jess, Alberto Bisin, and Matthew O. Jackson Volume 1. 1–30.
- Salas-Rojo, Pedro, and Juan Gabriel Rodríguez.** 2022. “Inheritances and Wealth Inequality: A Machine Learning Approach.” *The Journal of Economic Inequality* 20 (1): 27–51.
- Shapley, L. S.** 1953. “Stochastic Games.” *Proceedings of the National Academy of Sciences* 39 (10): 1095–1100.
- Shorrocks, Anthony F.** 1982. “Inequality decomposition by factor components.” *Econometrica: Journal of the Econometric Society* 193–211.

- Shorrocks, Anthony F.** 2013. “Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value.” *The Journal of Economic Inequality* 11 (1): 99–126.
- Solon, Gary.** 1999. “Chapter 29 - Intergenerational Mobility in the Labor Market.” In *Handbook of Labor Economics*, edited by Ashenfelter, Orley C., and David Card Volume 3. 1761–1800.
- Van Elk, Roel Adriaan, Egbert Jongen, Patrick Koot, and Alice Zulkarnain.** 2024. “Intergenerational Mobility of Immigrants in the Netherlands.” IZA Discussion Paper No. 17035.
- Vosters, Kelly.** 2018. “Is the Simple Law of Mobility Really a Law? Testing Clark’s Hypothesis.” *The Economic Journal* 128 (612): F404–F421.
- Vosters, Kelly, and Martin Nybom.** 2017. “Intergenerational Persistence in Latent Socioeconomic Status: Evidence from Sweden and the United States.” *Journal of Labor Economics* 35 (3): 869–901.

Appendix A: supplementary results

Table A1: Descriptive Statistics for the Income Analysis Sample

	Mean	SD	Mean	SD	% missing
Characteristics children					
Male	0.51	0.50			0
Family size	2.7	1.3			0
Household income	104664	73396			0
Second generation migrant	0.15	0.36			0
Third generation migrant	0.06	0.23			0
Family characteristics: measured at the household level					
Household income rank	0.50	0.29			0.005
Primary income share	0.80	0.39			0.006
Highest education	12.90	3.65			0.352
Total wealth rank	0.50	0.286			0.008
Bank and savings balances	52248	180882			0.008
Bonds and shares	36700	347042			0.008
House value	309723	379853			0.008
Entrepreneurial assets	15022	132238			0.008
Other real estate	30280	278106			0.008
Substantial interest	65576	1234938			0.008
Other assets	6103	112678			0.008
Total debt	159252	374265			0.008
Mortgage debt	134717	190709			0.008
Relationship status of household head(s) of child at age 15:					
Registered partners	0.82	0.38			0.023
Non-registered partners	0.04	0.19			0.023
Single parent	0.13	0.33			0.023
Other	0.01	0.11			0.023
Other family characteristics					
	Father		Mother		
Personal income	67636	49548	28588	20798	0.092
Personal earnings	82618	59506	32288	25831	0.155
<i>Most important source of income</i>					
Employment	0.67	0.41	0.54	0.43	0.042
Bonds or shares	0.04	0.17	0.01	0.09	0.042
Entrepreneurship	0.12	0.29	0.07	0.22	0.042
Substantial interest	0.00	0.05	0.03	0.12	0.042
Unemployment benefits	0.02	0.08	0.015	0.055	0.042
Welfare benefits	0.02	0.13	0.05	0.19	0.042

Other social security	0.00	0.05	0.01	0.06	0.042
Disability insurance transfers	0.08	0.23	0.06	0.20	0.042
Pension	0.02	0.10	0.04	0.14	0.042
Other	0.01	0.08	0.19	0.34	0.042
<i>Type of housing</i>					
Own house	0.74	0.41	0.70	0.42	0.052
Rental	0.06	0.19	0.106	0.257	0.052
Subsidized rental	0.21	0.36	0.20	0.33	0.052
Year of education	12.76	3.81	11.89	3.69	0.525
Average hourly wage	32.00	26.91	20.69	18.09	0.315
Most important sector of employment		In 68 categories			0.315
Suspected of any crime	0.07	0.25	0.02	0.15	0.013
Suspected of property crime	0.01	0.12	0.01	0.09	0.013
Suspected of violent crime	0.03	0.16	0.01	0.08	0.013
Suspected of other crime	0.04	0.20	0.01	0.11	0.013
Total health costs	2768	6589	2729	7769	0.013
General practitioner costs	178	139	201	148	0.061
Mental health care costs	234	2977	318	3293	0.061
Hospital care costs	1917	6430	1751	4697	0.061
Pharmaceutical care costs	538	2198	549	2035	0.061
Dental care costs	49	268	47	265	0.061
Age at first birth	29.7	4.9	27.0	4.4	0.018
Family size	4.2	2.4	4.1	2.3	0.214
Birth order	2.5	1.8	2.5	1.8	0.214
Father/mother not identified	0.02	0.15	0.00	0.05	0
Father/mother dead	0.01	0.09	0.00	0.07	0.018
Father/mother present in household	0.86	0.35	0.96	0.19	0.036
Country of origin father/mother		In 8 categories			0.018
Country of origin grandfather(s)		In 8 categories			0.392
Country of origin grandmother(s)		In 8 categories			0.328
<i>Extended family outcomes</i>					
Average income rank	0.50	0.22	0.50	0.22	0.237
Average education	12.57	3.16	12.70	3.10	0.414
Average wealth rank	0.53	0.23	0.53	0.23	0.236
Average health expenditures	2803	5157	2642	4943	0.228
% of siblings suspected of any crime	0.04	0.14	0.05	0.15	0.228

Note: this Table presents descriptive statistics of the income sample. The sample comprises of all 1,705,532 children born between 1980 and 1989 with non-missing income (96.7%). A detailed explanation of the variables can be found below this table.

Income. The construction of children’s and parents’ household income ranks is discussed in the main text.

The share of primary income represents the fraction of household income derived from labor, entrepreneurship, or capital. It is constructed similarly to parental household income. Specifically, for each parent, I calculate the primary income share for each year up to age 60—the same years in which household income is measured. The lifetime primary income share is then defined as the average of these yearly shares. Finally, the household share of primary income is determined by averaging the lifetime primary income shares of both parents.

Personal income refers to an individual’s income from labor, entrepreneurship, or transfers, measured at the personal rather than household level. As a result, it excludes partners’ incomes but also household-level income streams, such as capital gains or rental allowances. Personal earnings equals personal income minus income transfers. Following the same approach as before, I exclude years with income or earnings observations lower than €1000, and proxy a parent’s lifetime personal income and earnings by averaging all personal income and earnings observations up to age 60. Although the table above shows personal income and earnings in absolute values, in the analysis, I use ranks instead. The ranks are taken relative to all other parents in the sample.

In addition, I identify the primary sources of personal income, classified into 10 categories.⁴¹ Drawing on all yearly observations used in constructing the lifetime personal income measure, I first compute the most important source of income in each of those years. I then compute the fraction of years in which each category served as the main source of income.

Similarly, for each of those years, I calculate the fraction of years that the father or the mother lived in a self-owned house, a rental property, or a government-subsidized rental.

Wealth. The wealth variables are constructed in a manner analogous to the parental household income variable, as both are measured at the household level. I observe the values for each type of asset or liability of each parent in 2006. For each child, I determine the mean of the father’s and mother’s values for each asset or liability type.

The assets and liabilities included in this analysis are defined as follows. Bank and savings balances represent the total deposits held by a household in (savings) bank accounts, including foreign accounts. House value captures the market value of a household-owned dwelling used as the primary residence, while other real estate encompasses the total value of any additional properties owned by the household. Bonds and shares measure the combined value of bond and equity holdings, excluding ‘substantial interests’ (holdings of at least 5 percent of a company’s issued share capital), which are accounted for separately under the “substantial interests” variable. Entrepreneurial assets reflect the net balance of a household’s business-related assets and liabilities, and other assets include any remaining assets not covered by the aforementioned categories. Mortgage debt refers to debts associated with the household’s owner-occupied home, whereas other debt encompasses all other types of liabilities.

⁴¹One category is income from substantial interest. A substantial interest refers to a shareholder owning at least 5% of a company’s shares. This threshold is used for tax and regulatory purposes to identify large or influential shareholders. Income and wealth from such shares are measured separately.

Education. The education register reports individuals' highest completed level of education. I use this register to construct a years of education measure. Table A1 indicates that parental education information is absent for about 50 percent of the sample. This gap exists because Statistics Netherlands initiated systematic education data collection only in the late 1980s. Prior educational records are mainly sourced from large-scale surveys frequently administered by Statistics Netherlands and are also obtained indirectly from other government bodies, including the unemployment agency.

Occupation. I use monthly data on all employment contracts in the Netherlands from 2006 to 2009, collected by the tax authorities through third-party reporting. For each individual, I aggregate the total hours worked at each firm during this period. I then identify the firm where the individual has accumulated the most hours and assign the individual's employment sector based on that firm's classification. Sector categorizations are determined by the authorities in accordance with collective labor agreements. There are 68 sector categories in total, which include categories such as 'education and sciences', 'government defense', 'chemical industry', 'financial services', 'restaurants and bars', 'retail', etc. The average hourly wage is calculated by dividing the individual's total gross salary over the period by the total number of hours worked.

Healthcare. The health care expenditures are based on annual healthcare costs for care covered by the basic insurance. The basic insurance is legally mandated under the Healthcare Insurance Act for nearly all residents of the Netherlands. The costs refer to expenses for all types of care that are reimbursed by health insurers, and may include amounts ultimately paid by the insured themselves due to the deductible, but exclude copayments. If the insured received a bill and did not submit it to the insurer—e.g., because the deductible had not been reached—these costs are not included in the figures. The health care expenditures variables above are based on the subcategories of healthcare spending defined by Statistics Netherlands. For each of the subcategories, the annual costs are averaged over the period 2009 to 2011.

Crime. The crime data contains all offenses reported to the police since 2005. The data contain the reporting date, the offense type, and the individual identifier of the suspected offender(s) whenever there is a known suspect. I use these data to construct indicators of whether the father or the mother has been suspected of different types of crimes between 2005 and 2010.

Family structure. I record the family size and birth order of both the father and the mother by linking them to their siblings, which requires accessing the grandparents' identifiers. Consequently, these variables, along with any extended family outcomes, are missing for children whose grandparents cannot be identified. Additionally, I determine whether the father or mother was registered in the same household as the child at age 15 and classify the child's household type at that age into one of three categories: a couple with a registered partnership, a couple without a registered partnership, or a single-parent household. Furthermore, I calculate the parents' age at the birth of their first child and indicate whether either the father or the mother is not identified, as not all children have both parents identified.

Migration background. I have information on the country of origin of all identified parents and grandparents. I distinguish eight regions: the Netherlands, Morocco, Turkey, Surinam, Dutch Antilles, Western Europe, Eastern Europe, and others.

Extended family outcomes. For each parent separately, I determine the mean years of education, household income rank, wealth rank, and annual health expenditures across all their siblings. Additionally, I calculate the fraction of these siblings who have been suspected of committing a crime.

Table A2: Family Background Characteristics across the Predicted Income Distribution

	<i>Predicted Income Bins</i>								
	0- 0.5	0.5- 1	1- 5	5- 10	10- 90	90- 95	95- 99	99- 99.5	99.5- 100
Child income rank	16.89	19.45	25.2	31.11	50.54	65.78	70.41	74.63	78.12
	<i>Family background characteristics</i>								
Parental income rank	5.45	7.42	11.47	16.33	49.22	87.74	93.55	97.15	98.47
Parental wealth rank	11.75	12.5	14.11	17.26	50.81	74.38	80.2	86.55	90.06
Max. education parents	7.94	8.79	9.43	9.86	13.04	16.08	16.68	17.08	17.26
Health costs parents	5508	5096	4539	3971	2681	1929	1907	1899	1941
Crime father	0.58	0.45	0.3	0.18	0.05	0.02	0.02	0.03	0.03
Father presence	0.33	0.35	0.46	0.62	0.88	0.97	0.98	0.98	0.98
Migration background	0.3	0.34	0.48	0.52	0.18	0.1	0.12	0.13	0.16
Age at first birth mother	22.1	22.9	24.1	25.3	27.1	28.4	28.7	29.0	28.9
Extended family income	16.69	20.07	25.06	30.29	49.22	64.71	69.53	75.2	79.61
Extended family wealth	22.51	24.53	27.84	32.14	52.64	65.55	68.96	72.62	76.29
N	1,705	1,706	13,644	17,055	272,886	17,055	13,644	1,706	1,706

Notes: each column shows descriptive statistics for a group of children in the test data from the same predicted income bin. The predicted income bins are constructed by predicting the income ranks of all children in the test data using the model with all explanatory variables, ranking them from low to high, and sorting them into bins according to their position in the predicted income distribution. All values are averages, with missing values excluded from the calculations. Health expenditures parents equals the average health expenditures of the father and mother between 2009 and 2011. Extended family income (wealth) is calculated as the average income (wealth) rank of the father’s and mother’s siblings. Migration background is an indicator which equals 1 if the child is a second or third generation migrant. The other variables are discussed in Table 1.

Table A3: Predicting Child Income using Smaller Samples

Share of core sample	Test data sample size	R^2	Mean income of bottom 0.5%	Mean income of top 0.5%
(1)	(2)	(3)	(4)	(5)
0.01	3,411	0.125	0.30	0.61
0.02	6,822	0.143	0.26	0.81
0.05	17,056	0.153	0.23	0.71
0.1	34,111	0.157	0.20	0.72
0.2	68,222	0.162	0.19	0.75
0.4	136,443	0.165	0.17	0.77
0.6	204,664	0.168	0.18	0.77
0.8	272,885	0.168	0.17	0.77

Notes: this Table presents estimates of explanatory power for gradient-boosted decision trees that include all explanatory variables (as in Figure 2), using smaller samples. Column 1 reports the share of the core sample that is used for the analysis. Column 2 reports the sample size of the test-data. Columns 3 reports the R^2 . Columns 4 and 5 report the mean observed income of the 0.5 percent of children with the lowest and highest predicted income, respectively. Each model is trained on a randomly selected 80% of the respective sample, and evaluated on the remaining 20%.

Table A4: Predicting Child Income: Varying Years and Ages of Income Measurement

	R^2	Mean income of bottom 0.5%	Mean income of top 0.5%
Years of income	A. Varying years of income measurement		
1	0.139	0.17	0.76
2	0.145	0.18	0.78
3	0.153	0.18	0.77
4	0.154	0.18	0.78
5	0.159	0.17	0.78
6	0.159	0.17	0.79
7	0.163	0.18	0.79
8	0.163	0.18	0.78
9	0.167	0.18	0.79
All	0.174	0.17	0.79
All > 32	0.171	0.17	0.79
Age child	B. Varying ages of income measurement		
30-33	0.130	0.18	0.77
34-37	0.156	0.18	0.78
38-41	0.158	0.18	0.79

Notes: each row presents the R^2 and corresponding the mean observed income of the 0.5 percent of children with the lowest and highest predicted income based on gradient-boosted decision trees that include all explanatory variables to predict child income (as in Section 4). The analysis sample consists of all 330,018 children born in 1980 and 1981 for whom I observe all incomes between ages 30 and 41. Each model is trained on the same randomly selected 80% of this sample, and evaluated on the remaining 20%. Panel A varies the number of years of income data used to construct the child income rank. The one-but-last row in panel A uses all income observations, as in the main results. The last row uses all income data above age 32. Panel B uses four years of income data, but varies the ages at which income is measured.

Table A5: Within Neighborhood Estimates

	Household income rank	
	(1)	(2)
Predicted income rank	1.001 (0.002)	0.965 (0.002)
Neighborhood Fixed Effects		x
N	1,664,192	1,664,192
R^2	16.9%	17.4%

Notes: columns 1 and 2 report results from separate regressions of a child's income rank on its predicted value, with and without neighborhood fixed effects. The sample corresponds to all children from the core sample with an available neighborhood identifier at age 15. The predictions are generated in the following steps: (i) randomly split this sample into five folds, (ii) leave out one fold and estimate a gradient boosted decision tree with the same tuning parameters as the model in Figure 2 on the remaining folds, (iii) generate predictions for all observations in the omitted fold, and (iv) repeat this step until all folds are held out once. This procedure ensures that predictions are always made out of sample. Standard errors, shown in parentheses, are clustered at the neighborhood level in column 2.

Table A6: Alternative Neighborhood Intergenerational Mobility Estimates

A. Main results	
Signal SD(AUM)	0.044
Signal SD(MAUM)	0.020
B. Estimation with a second order polynomial	
Signal SD(AUM)	0.047
Signal SD(MAUM)	0.021
C. Estimation with a third order polynomial	
Signal SD(AUM)	0.048
Signal SD(MAUM)	0.020
D. Estimation at the 5th percentile of the parental/predicted income distribution	
Signal SD(AUM)	0.050
Signal SD(MAUM)	0.023
E. Estimation at the 75th percentile of the parental/predicted income distribution	
Signal SD(AUM)	0.032
Signal SD(MAUM)	0.019
F. Relative intergenerational mobility	
Signal SD($\hat{\beta}_n$)	0.150
Signal SD($\hat{\gamma}_n$)	0.071
G. (Multidimensional) Absolute Upward Mobility estimates at the municipality level	
Signal SD(AUM)	0.034
Signal SD(MAUM)	0.016

Notes: each row reports the signal standard deviation for a distinct set of estimates. Panel A restates the results in Figure 5. Panel B and Panel C report results for AUM and MAUM when equations 5 and 6 are estimated with second and third order polynomials. Panel D reports results when the fifth percentile of the national parental or predicted income distribution is used in equation 5 or 6. Panel E reports results when the 75th percentile is used. Panel F reports the signal standard deviation of the slope coefficients in equations 5 and 6. Panel G reports estimates at the municipality level. Estimates in Panel A, B, C, D, and G use weights based on the number of children with below median income in each neighborhood or municipality. Estimates in Panel E are weighted by the number of children with above median parental income. Estimates in Panel F are weighted by the number of children in each neighborhood. Signal standard deviation estimates are computed by subtracting the weighted average squared standard error from the weighted sample variance of the estimates, then taking the square root.

Table A7: The Relationship between Neighborhood Upward Mobility and Predicted Income

	Predicted income
Absolute Upward Mobility	0.674 (0.004)
N	166,307

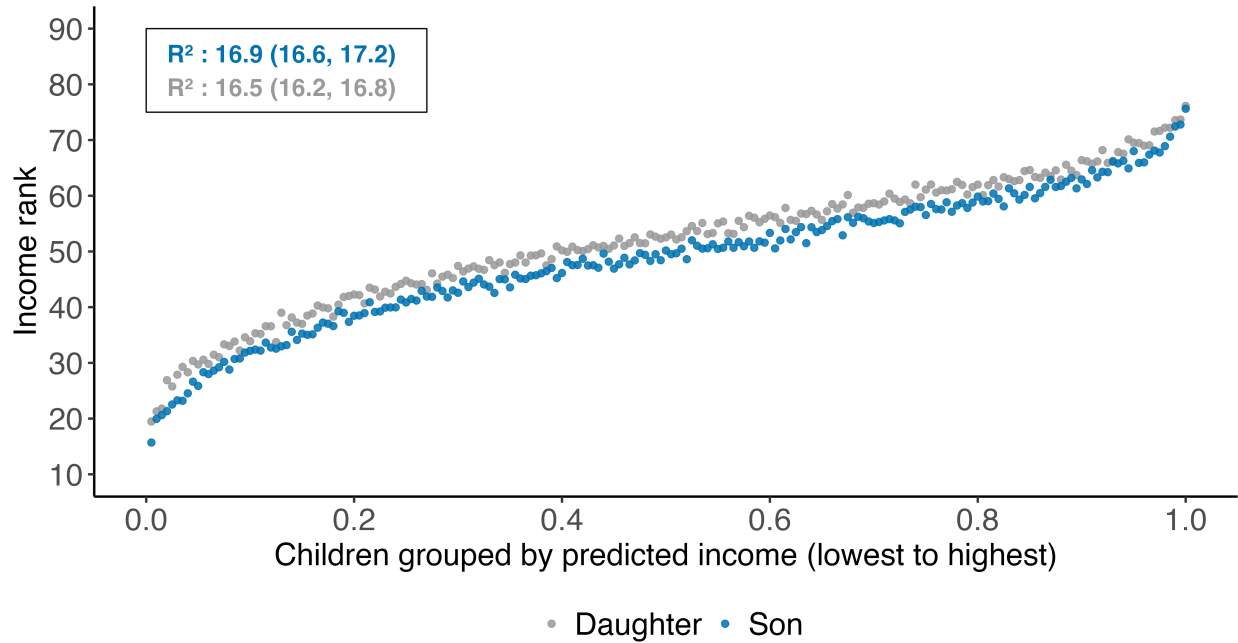
Notes: this Table reports the results from a regression of predicted income on neighborhood upward mobility. Parental household income is included as a control variable. The sample includes all children from the core analysis sample with a parental income rank between 20 and 30. The income predictions are generated by gradient-boosted decision trees with as explanatory variables all family characteristics, as further explained in Table A5. The Absolute Upward Mobility (AUM) variable corresponds to the AUM estimate of the neighborhood where a child was registered at age 15 (see Section 5).

Table A8: Relaxing the single-index restriction in MAUM

	A. Results across all 2,829 neighborhoods
Fraction of rejected F-tests	0.100
SD in absolute upward mobility	0.036
Signal SD in absolute upward mobility	0.022
	B. Results across all 355 municipalities
Fraction of rejected F-tests	0.194
SD in absolute upward mobility	0.017
Signal SD in absolute upward mobility	0.016


Notes: This table reports the robustness tests discussed in Section 5. Panel A is constructed in three steps. For each neighborhood, I (i) estimate the restricted (Equation 6) and unrestricted (Equation 7) regressions; (ii) conduct an F-test of the restricted specification against the unrestricted; and (iii) compute a new measure of neighborhood absolute upward mobility by inserting $\hat{\kappa}_n$ and $\hat{\gamma}_{n,1}, \dots, \hat{\gamma}_{n,9}$ into Equation 7, fixing each index at its average value among children at the 25th percentile of $\hat{Y}(\mathbf{X}_f)$. Panel A reports the share of neighborhoods in which the F-test rejects at the 5 percent level, together with the raw and signal standard deviations of the new mobility measure. Panel B applies the same procedure at the municipality level.

Figure A1: Predicting Children's Household Income Rank by Gender



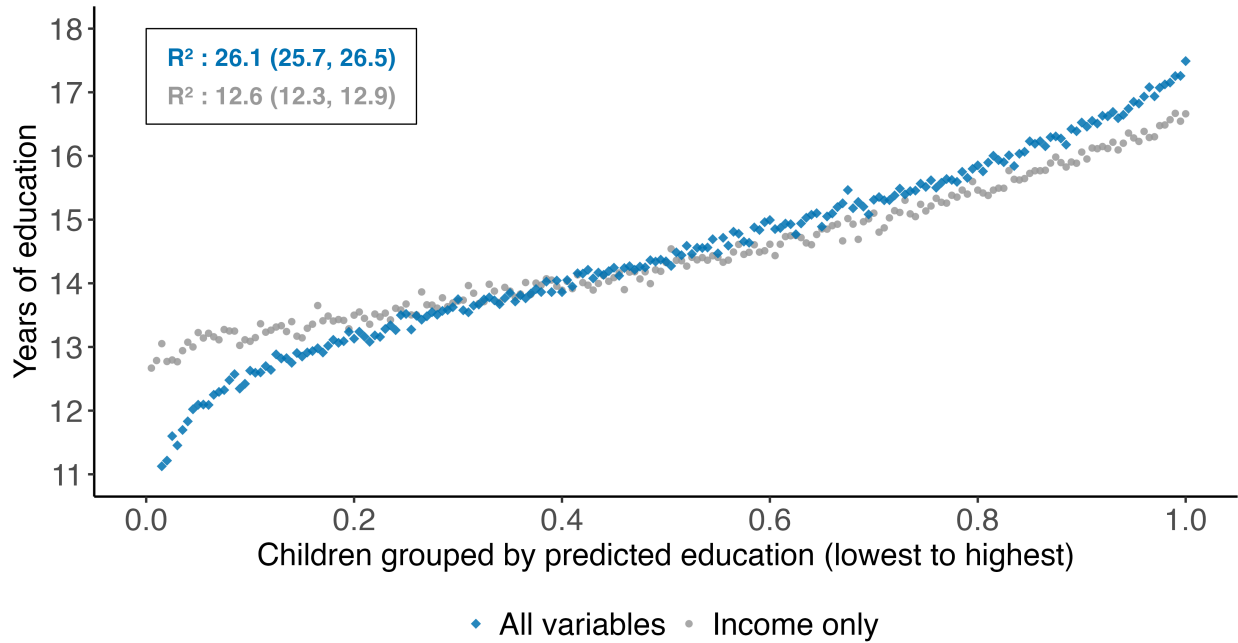
Notes: this Figure presents binscatter plots of sons' and daughters' household income ranks for 173,948 sons and 167,159 daughters in the test data, who are sorted into bins based on their predicted income rank. Predictions are generated using the same predictive model and explanatory variables as in Section 4, now applied separately to each gender. The construction of the graphs follows the same steps as in Figure 2, now separately for each gender. Confidence intervals for the R^2 are bootstrapped from the test data using 599 samples and are reported in brackets

Figure A2: Predicting Children's Personal Income by Gender




Notes: this Figure presents binscatter plots of sons' and daughters' personal income ranks for 172,976 sons and 164,990 daughters in the test data, who are sorted into bins based on their predicted income rank. The graphs are constructed using the same steps as in Figure 2, applied to children's personal income ranks instead of household income ranks. Confidence intervals for the R^2 are bootstrapped from the test data using 599 samples and are reported in brackets

Figure A3: Predicting Children's Education and Crime




(a) Education



(b) Crime (sons only)

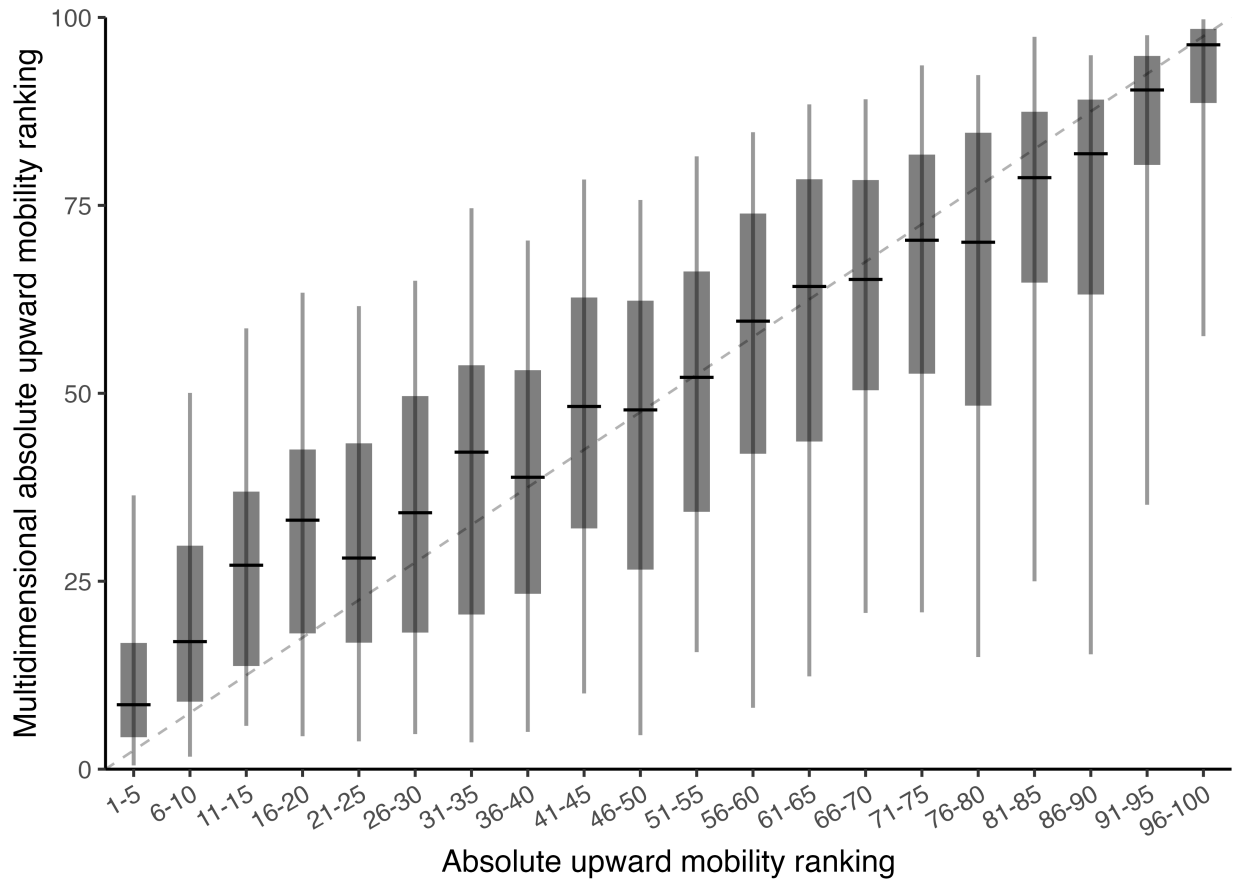
Notes: the Figures above present binscatter plots of children's years of education and crime for two predictive models. The crime outcome is an indicator of whether a child has been suspected of any violent crime between ages 20 to 33. The children are sorted in 200 bins from lowest (0) to highest (1) predicted education/crime. The education analysis only includes children born between 1985 and 1989. The crime analysis only includes sons born between 1985 and 1989. Panel A reports results for 174,958 children from the test sample. Panel B reports the results for 89,176 sons from the test sample. The blue and grey graphs are constructed using the same steps as in Figure 2. Confidence intervals for the R2 are bootstrapped from the test data using 599 samples and are reported in brackets

Figure A4: The 30 Most Predictive Family Characteristics of Child Income



Notes: each row displays a boxplot representing the distribution of Shapley values for a given variable. These Shapley values are computed using the algorithm of Lundberg et al. (2020) for each variable and each child from the test dataset. The variables shown are those with the 30 highest mean absolute Shapley values across these observations. The whiskers indicate the 2.5th and 97.5th percentiles, the box edges correspond to the 25th and 75th percentiles, and the center bar represents the median.

Figure A5: Comparing Neighborhood Upward Mobility Rankings



Notes: this Figure compares neighborhood rankings based on Absolute Upward Mobility with rankings based on Multidimensional Absolute Upward Mobility. All 2,829 neighborhoods are ranked from 0 to 100 on each measure based on the estimates in Figure 5. Neighborhoods are grouped into twenty equally sized bins on the horizontal axis according to their Absolute Upward Mobility ranking. The histograms show the distribution of the Multidimensional Absolute Upward Mobility rankings among neighborhoods in the same bin. The outer lines span the 5th and 95th percentiles. The inner boxes span the 25th and 75th percentiles. The horizontal segments mark the medians.

Appendix B: intergenerational mobility estimates

Additional results. Given that my baseline intergenerational mobility estimate differs from other estimates in the Netherlands, I provide additional estimates here that are commonly reported in the literature. These can be used by other researchers who wish to make cross-country comparisons. Below, I also present a sensitivity analysis and elaborate on why my estimates differ from prior estimates.

Table B1 reports the rank-rank correlation as well as the Intergenerational Income Elasticity (IGE) using logs of household income instead of ranks in columns 1 and 2. These are both close to 0.33. Columns 3 and 4 report results for sons and daughters separately and rely on children’s personal income ranks instead of household income ranks. These estimates are very similar between genders and somewhat lower than the rank-rank correlation based on household income. Finally, column 5 reports the sibling correlation in income, which equals the adjusted R^2 of a regression of child income on sibling fixed effects. This estimate suggests that about 30% of all variation in income ranks is driven by factors shared between siblings.

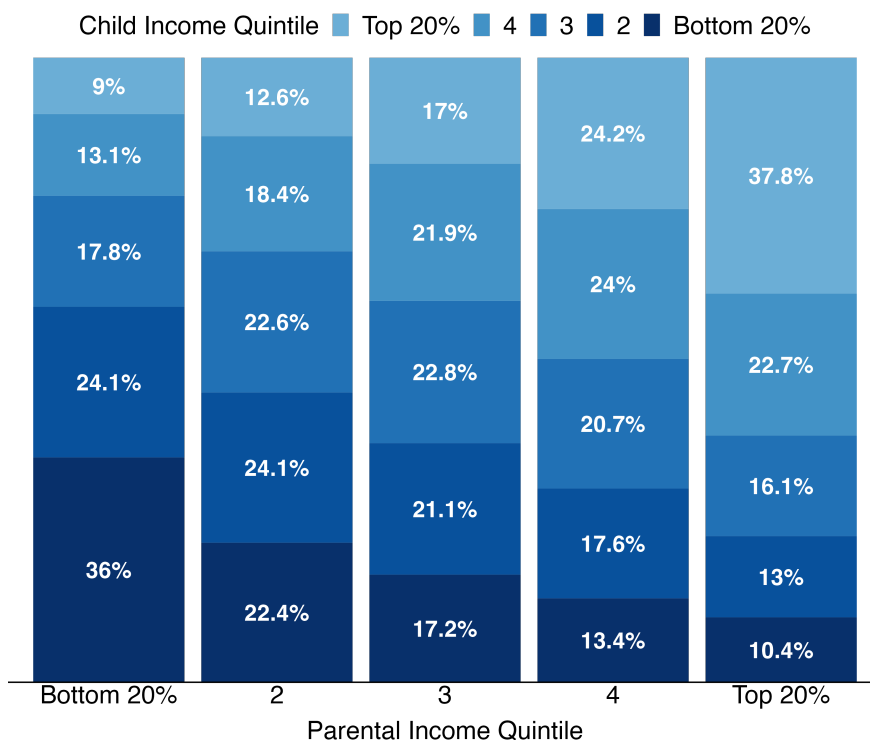
Figure B1 reports a 5×5 transition matrix. This table can be used to compare upward or downward mobility estimates across countries.

Table B1: Intergenerational mobility estimates

	Rank rank correlation	IGE	Personal income rank (daughters)	Personal income rank (sons)	Sibling correlation
	(1)	(2)	(3)	(4)	(5)
Coefficient	0.329 (0.001)	0.341 (0.001)	0.293 (0.001)	0.297 (0.001)	-
N	1,712,291	1,712,291	871,371	833,931	1,705,532
Adjusted R^2	0.108	0.095	0.097	0.100	0.301

Notes: column (1) shows results from a regression of a child’s household income rank on the parents’ household income rank. Column (2) shows results from a regression of the log of child household income on the log of parental household income. Columns (3) and (4) show results from a regression of sons’ or daughters’ personal income rank on parents’ household income rank. Column (5) reports the sibling correlation. This is estimated by the adjusted R^2 of a regression of child income on sibling fixed effects. The sample includes the core analysis sample (Table A1) excluding observations with missing parental income (0.9 percent). Standard errors are in parentheses.

Figure B1: Transition matrix



Notes: this Figure presents the transition matrix of child income conditional on parental income quintile. Each bar represents the distribution of child income quintiles for children whose parents fall in the corresponding parental income quintile on the X-axis. The segments within each bar show the share of children reaching each income quintile, as indicated by the color legend. The sample ($N = 1,702,355$) includes the core analysis sample (Table A1) excluding observations with missing parental income (0.9 percent).

Sensitivity. Next, I evaluate the sensitivity of the rank-rank correlation of 0.33 to various specification choices. Although it would be ideal to perform robustness checks using the full analysis sample, the specific data requirements for each check necessitate the use of different samples. Stability of the estimates within these samples strengthens confidence that the estimates would also remain stable under different specifications in the broader analysis sample.

Table B2 reports mobility estimates using varying years of income information of parents. I focus on all children for whom both the father and the mother have at least 9 observable income observations. The estimates attenuate somewhat with fewer years of income, but the change in the rank-rank correlation is limited after 5 years of income are used. This suggests that attenuation bias is unlikely to be an issue.

Table B3 reports mobility estimates using incomes of parents measured in different periods. I focus on all children for whom parental income is observed between 2001 and 2011. I average income over 4 years for each of the specifications. The estimates are very similar, regardless of when parental income is measured.

Table B4 reports mobility estimates using incomes of children measured at varying ages. I focus on all children born in 1980 or 1981 for whom all incomes are observed between ages 30 to 41. I average income over 5 years for each of the specifications. The estimates

show that measuring income early attenuates the estimates, but they stabilize after age 34. Overall, the differences are relatively small.

Table B2: Intergenerational mobility estimates: varying years of parental income

Years of income	1	2	3	4	5	6	7	8	9
Coefficient	0.296 (0.001)	0.307 (0.001)	0.313 (0.001)	0.317 (0.001)	0.321 (0.001)	0.324 (0.001)	0.325 (0.001)	0.326 (0.001)	0.327 (0.001)
N	1,250,211	1,250,211	1,250,211	1,250,211	1,250,211	1,250,211	1,250,211	1,250,211	1,250,211
R^2	0.082	0.088	0.091	0.094	0.096	0.098	0.098	0.099	0.099

Notes: each column presents results from a regression of a child's household income rank on the parents' household income rank. The number of years of income data used to construct the parental income rank varies across columns, as indicated in the first row. The income observations used are always those closest to age 35. Standard errors are reported in parentheses. The sample consists of all children for whom at least 9 paternal *and* 9 maternal incomes are available.

Table B3: Intergenerational mobility estimates: measuring parent income at different ages

	(1)	(2)	(3)
Coefficient	0.295 (0.001)	0.296 (0.001)	0.285 (0.001)
Years of income measurement parents	2001-2004	2004-2007	2007-2011
N	1,340,585	1,340,585	1,340,585

Notes: each column presents results from a regression of a child's household income rank on the parents' household income rank. Child income ranks are measured as in the main analysis in this paper. Parent household income ranks are always based on 5 years of income, but the periods at which incomes are measured vary across columns. The sample consists of all children in the core sample for whom parental income is observed between 2001 and 2013. Standard errors are reported in parentheses.

Table B4: Intergenerational mobility estimates: measuring child income at different ages

	(1)	(2)	(3)
Coefficient	0.277	0.308	0.314
	(0.002)	(0.002)	(0.002)
Age child	30-33	34-37	38-41
N	329,943	329,943	329,943

Notes: each column presents results from a regression of a child’s household income rank on the parents’ household income rank. Parent household income is measured as in the main results of this paper. Child household income ranks are always based on 4 years of income, but the ages at which child incomes are measured vary across columns. The sample consists of all children for whom all incomes between ages 30 and 41 are available. Standard errors are reported in parentheses.

Comparison with other studies. There are three recent estimates of the rank-rank correlation in the Netherlands.

Most closely related is Van Elk et al. (2024). They study intergenerational mobility differences among migrants and natives, and use the same data as in this paper. While in the main paper they focus on disposable household income, in the Appendix, they report a rank-rank correlation of 0.22 that corresponds to gross household income. There are four main differences between our approaches. Below, I describe these differences and quantify their importance in Table B5 step by step.

The core analysis sample in this paper includes all children born between 1980 and 1989, excluding only 3.3% of children with missing income observations. Van Elk et al. consider children born between 1983 and 1988. Column 1 of Table B5 replicates the rank-rank correlation for children born in these years. For these cohorts, I find a similarly large rank-rank correlation of 0.33. Starting from this baseline estimate, I change my measurement approach so as to align with Van Elk et al.

First, Van Elk et al. drop all children who do not live independently in 2003 and who do live independently in 2017 to 2019, whereas I do not make such sample restrictions. Dropping these individuals results in a 23 percent smaller sample and reduces the rank-rank correlation by 0.021 (columns 2 and 3).

Second, Van Elk et al. measure child income from 2017 to 2019, when children are aged 29 to 36. I average income over all available observations from age 30 onward and up to 2024. Implementing their age at measurement further reduces the rank-rank correlation by 0.025 (column 4).

Third, Van Elk et al. measure parental income from 2003 to 2005. I measure parental income over all available observations from 2001 and up to age 60. On average, that corresponds to 13 observations for fathers and 16 observations for mothers. Implementing their parental age at measurement further reduces the estimate by 0.026 (column 5).

Fourth, Van Elk et al. define parents as the head of the child’s household in 2003 and his or her partner. Parental income is then defined as the income of this household head and his or her partner between 2003 and 2005. Instead, I define parents based on legal relationships, as documented in the ‘parent-child register’. Following Chetty et al. (2014), parental income is then defined as the average of the household income of the father and the

mother. Our parental income concepts align when the child, father, and mother live together between 2003 and 2005. However, when at least one of the legal parents is not present in the household in these years, our definitions differ. Implementing their measure of parental income reduces the estimate by 0.060, resulting in an estimate that is very close to their main estimate (column 6).

This drop is relatively large because the legal father or mother is absent from the child’s household in 2003 in 28 percent of cases. For these children, the income of the legal parents is considerably more predictive than that of their household heads.⁴²

Table B5: Comparison with Van Elk et al. (2024)

	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient	0.334	0.319	0.313	0.288	0.262	0.202
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	1,020,416	875,285	771,842	771,842	771,534	771,534
Adjustments						
Child born between 1983-1988	x	x	x	x	x	x
Child not living independently in 2003		x	x	x	x	x
Child living independently in 2017 to 2019			x	x	x	x
Child income measured in 2017 to 2019				x	x	x
Parental income measured in 2003 to 2005					x	x
Parental income based on household head						x

Notes: each column presents an estimate of the rank-rank correlation, using different variable definitions and sample selections. The specification in Column 1 uses the same variable definitions as in the main text (Table B1), focusing exclusively on children born between 1983 and 1988. The subsequent columns report results using a different sample selection or different variable definitions. These differences are further explained in the main text above.

There are also two estimates of the rank-rank correlation which are based on different data. First, Boustan et al. (2025) compare intergenerational mobility among migrants and natives in 15 destination countries, including the Netherlands. While the children’s incomes are based on the same population-wide administrative data, the parents’ incomes in their study are based on a random sample of administrative data from before 2001 (in Dutch: the ‘IPO’). This random sample contains incomes in 1981, 1985, and annually from 1989 for about 3.3% of the population. Boustan et al. report intergenerational mobility estimates of

⁴²Using the sample of children for whom at least one of the legal parents is not present in the household, I find a rank-rank correlation of 0.32 when using the legal parents’ incomes. This drops to 0.05 when using the household head and his/her partner’s income.

0.24 and 0.22 for sons and daughters born between 1982 and 1987 (See Table C.9.23). There are three main differences with their approach: (i) they use children’s personal income (in 2018 and 2019), whereas I use household income (measured above age 30 and up to 2023), (ii) they use the sum of parents’ personal incomes instead of parental household income, and (iii) they measure parental income from 1998 to 2004, whereas I measure parental income from 2001 and up to age 60.

The personal income measure of Statistics Netherlands excludes not only the partner’s income but also income components from joint tax statements that cannot be attributed to specific individuals. These include income from wealth and allowances allocated based on household-level income, such as child and rental allowances. Consequently, the sum of parents’ personal incomes does not match the household income measure provided by Statistics Netherlands, which I employ in this study, even for cohabiting parents.

I do not have access to the survey, precluding a direct comparison with my results. However, in Table B6, I try to mimic their analysis as closely as possible, using the population wide administrative data. I begin by restricting my sample to children born between 1982 and 1987 and estimate the baseline rank–rank specification, which yields a correlation of 0.33.

In column 2, I revise the parental income measure to the sum of both parents’ personal incomes from 2001 to 2007. While I cannot observe incomes prior to 2001, this at least aligns the number of years over which parental income is measured.⁴³ This reduces the rank–rank correlation to 0.295. I then replace my original outcome with the child’s personal income rank, based on income measured in 2018 and 2019. This further reduces the estimate to 0.256 (column 3), which is quite close to their estimate. Remaining differences may reflect discrepancies between survey and administrative data, for instance due to missing income information for non-cohabiting parents in the survey.

Lastly, Manduca et al. (2024) study trends in absolute mobility across multiple countries. While their main goal is not to quantify relative intergenerational mobility, they also report rank-rank correlations for the Netherlands from 0.23 in the 1974 cohort to 0.16 for the 1984 cohort. They use a very similar approach as Boustan et al. They also link children’s incomes from the population wide administrative data to parental income from the representative survey, and also rely on personal income measures for children and parents. The main difference between Manduca et al. and Boustan et al. is that Manduca et al. measure parental and child income in only one year (the closest observation to age 30 for both generations). As shown in Table B6 column 4, using only one income observation for parents and children and measuring child income at age 30 further reduces the estimate to 0.23.

Since I do not observe parents’ incomes before 2001, I cannot assess the impact of also measuring parental income at age 30. However, Table B4 shows that results attenuate somewhat when measuring child incomes in the early 30s, suggesting that individuals may not be on their long-term income trajectory at that age. A similar bias may occur when measuring parental income at this relatively young age.

⁴³Estimates are stable across different years of income measurement between 2001 and 2011 (Table B3). This makes it likely that estimates are also similar when parental income is measured between 1998 and 2004.

Table B6: Comparison with Boustan et al. (2025) and Manduca et al. (2024).

	(1)	(2)	(3)	(4)
Coefficient	0.331	0.295	0.256	0.228
	(0.001)	(0.001)	(0.001)	(0.001)
N	992,192	992,192	992,192	992,192
Adjustments				
Child born between 1982 and 1987	x	x	x	x
Using personal income of parents		x	x	x
Using personal income of child in 2018 and 2019			x	x
Using one income observation for parents (in 2001) and children (at age 30)				x

Notes: each column presents an estimate of the rank-rank correlation, using different variable definitions and sample selections. The specification in Column 1 uses the same variable definitions as in the main text (Table B1), focusing exclusively on children born between 1982 and 1987. The subsequent columns report results using a different sample selection or different variable definitions. These differences are further explained in the main text above.

Appendix C: Supplementary Results with Adoptees

How much of the income gaps in Figure 2 are driven by pre-birth versus post-birth factors? Prior research used adopted children to decompose standard intergenerational mobility coefficients in these two factors (Sacerdote (2007), Holmlund et al. (2011), Fagereng et al. (2021)). This section extends that agenda by applying such a decomposition to a much broader, multidimensional measure of family advantage and by asking which dimensions of family background are rooted more strongly in each of the two components.

I consider a sample of 5,051 international adoptees born between 1980 and 1989 and who arrived in the Netherlands within six months of birth.⁴⁴ These children are not genetically related to their adoptive parents and were not cared for by them during pregnancy and shortly after birth, but have been raised by them since they were at most six months old. This unique context makes them an interesting group for studying the importance of the post-birth environment.

I begin by comparing adoptees' observed income to their predicted income based on all background characteristics of their adoptive families, using the same model as in Figure 2.⁴⁵ Figure C1 shows that adoptees consistently have lower average income than predicted. This aligns with earlier evidence that adopted children tend to perform worse in education and in the labor market (Sacerdote (2011)). Strikingly, the income of adopted children remains low even if they are raised in highly advantaged families. For example, adoptees with a predicted income rank around 70 have an average income rank of 44, despite being raised in families whose income, wealth, education, and other characteristics are among the highest in the population (Table C1).

The linear regression slope of 0.28 indicates that being raised in a family that is associated with a 1 rank higher income for own-birth children increases the income of adoptees by only 0.28 ranks. Under the assumptions that (i) there is no correlation between adoptees' genetic endowments or health at infancy and their adoptive family background characteristics, and (ii) the results can be generalized to the broader population, these estimates suggest that around 30 percent of the disparities in Figure 2 are shaped by the post-birth environment. The remaining share must under these assumptions reflect differences in pre-birth factors such as genetic endowments or prenatal conditions.⁴⁶

I do not have access to information on matching procedures from this period, restricting a comprehensive assessment of assumption 1. However, Table C2 shows that the estimate stays close to 0.28 even when I compare children of the same gender, who come from the same country, and who were adopted in the same year. This suggests that selection on these observable characteristics is at least of limited empirical importance.⁴⁷

⁴⁴Although the Netherlands lacks an adoption register, Statistics Netherlands developed a reliable method to identify adoptees. They sent a survey to a random subset of all plausible adoptees to verify their method. Overall, 97.8 percent of respondents in my sample confirmed they were adopted ($n = 787$).

⁴⁵All families with adopted children were excluded from the training data. As a result, all predictions in Figure C1 are out-of-sample.

⁴⁶The magnitude is in line with prior estimates. Sacerdote (2007) and Holmlund et al. (2011) attribute 20 to 30 percent of the intergenerational persistence in education to post-birth factors in the US and Sweden, and Fagereng et al. (2021) attribute roughly half of wealth persistence to post-birth factors in Norway.

⁴⁷The excess demand for infant adoptees in the 1980s likely discouraged selective placement, as prioritizing specific characteristics would have significantly increased already long waiting times. Waiting times during

To assess external validity, Figure C1 also reports results for 3,802 own-birth children from families with at least one adopted child. Their observed income aligns closely with the prediction, suggesting that intergenerational transmission for own-birth children in these families operates similarly as in the broader population. I also test for heterogeneity among adoptees from the six largest origin countries in Table C3 and do not reject equality of the coefficients ($p = 0.86$). The much weaker link between adoptees' income and the characteristics of their adoptive families thus appears a robust phenomenon across migration backgrounds.⁴⁸

Which family characteristics most strongly predict adoptees' income? Identifying which adoptive family background characteristics are most strongly associated with adoptees' income is challenging because nonparametric predictors like decision trees can perform poorly in smaller samples. Estimating a new gradient-boosted decision tree using the sample of international adoptees results in a negative R^2 . An OLS regression where all variables enter linearly results in a slightly higher adjusted R^2 of 1.4 percent, but it has many imprecisely estimated coefficients, rendering it difficult to interpret.

To make the analysis tractable, I instead collapse the full set of family background variables into the nine Shapley-value indices from Section 4. Each index represents a different family background dimension and measures its contribution to a child's predicted income. The wealth index, for example, is the total contribution of all wealth-related variables to a child's predicted income. By construction, the nine indices sum to the full prediction (Equation 4). The main advantage of this approach is that it collapses many correlated variables into a single measure for each category, while assigning greater importance to predictors that are more predictive of own birth children's income. The indices are standardized to have mean zero and variance one.

I then regress child income on the nine indices, separately for own-birth children and adoptees. Table C4 column 1 reports the results for own-birth children. The coefficient on the parental income index is 0.063, meaning that children in families who are one standard deviation 'more advantaged' in parental income, but similar in all other dimensions, have 6.3 ranks higher income on average. All coefficients carry the same sign because every index is constructed so that higher values correspond to more advantaged family backgrounds. A higher parental crime index, for example, reflects less parental crime, not more.

Consistent with the main results, the joint explanatory power of the indices is 17.2 percent. The strongest predictors are parental income, parental wealth, and extended family outcomes. These results illustrate that the Shapley value indices provide a simple way to collapse the full set of predictors into a small number of indices while preserving the explanatory power of the original model.

Column 2 reports the same regression for the sample of own-birth children in families with adopted children. These coefficients are very similar to the coefficients in column 1. This shows that also the partial relationships between specific family background dimensions and child income operates similarly in these families as in the broader population.

this period could span several years. See, for example, the government report 'Rapport Commissie Onderzoek Interlandelijke Adoptie' (in Dutch, 2021).

⁴⁸Extrapolation remains difficult when adoptive parents raise adopted children differently or when adopted children's initial endowments such as infant health differ. The quantitative importance of these concerns, which I share with the broader literature using adopted children, remain an open question.

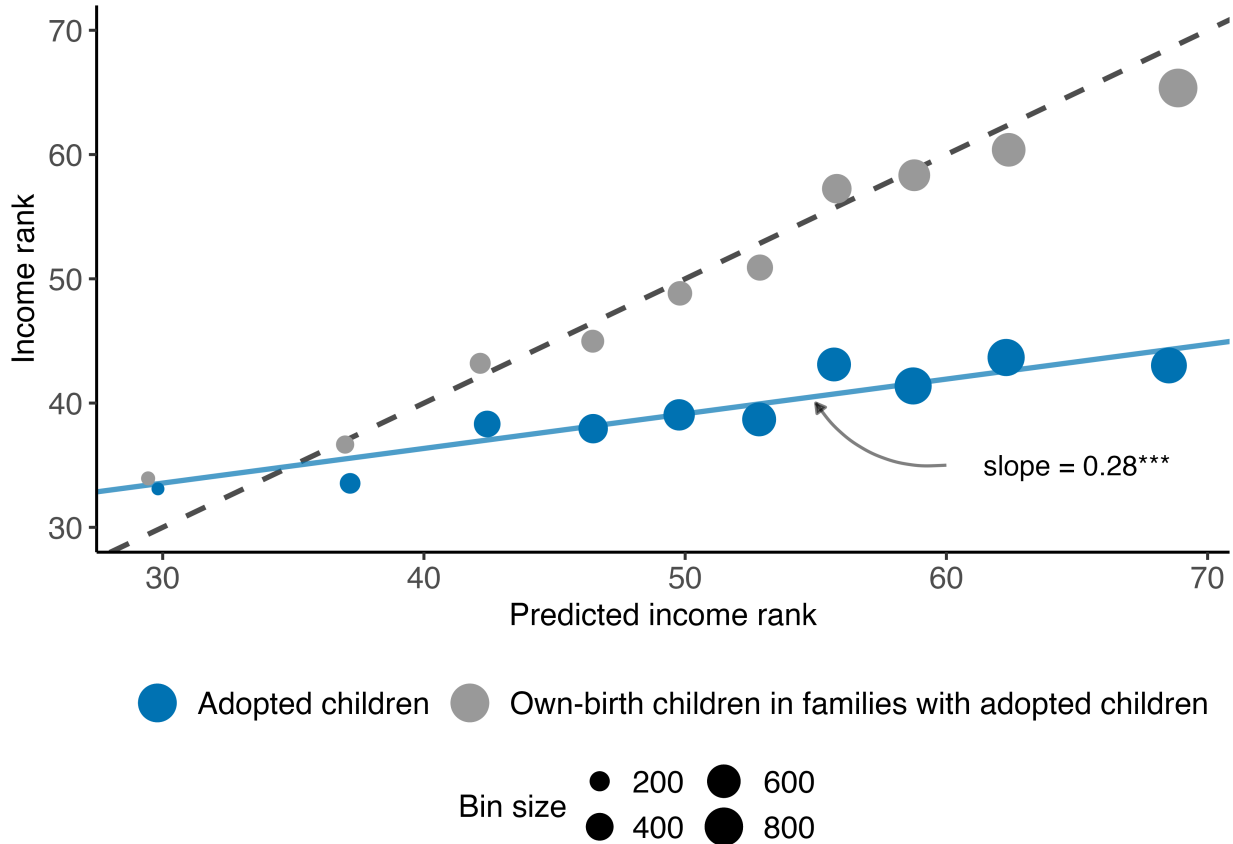
The coefficients are very different for the adopted children, reported in column 3. The associations are uniformly much weaker. The parental income coefficient, for example, is about one-fourth of its own-birth counterpart: a one standard deviation increase in the index raises adoptees' income by only 1.6 rank. The one exception is parental crime, which shows no attenuation, but the estimate is imprecise because crime is rare among adoptive parents.⁴⁹

The extent of attenuation differs across dimensions. It is strongest for parental income, extended family outcomes, health expenditures, and education. In fact, the wealth and family-structure variables now have comparable estimates as the parental income index. One interpretation is that parental income, education, and extended family outcomes are more closely tied to genetic endowments that parents transmit to their biological children. These channels are absent for adoptees, so the associations with these indices falls sharply once the biological link is removed. Wealth, occupation, and family structure, in contrast, may operate more through the environment in which children are raised, and so retain more of their predictive power for adoptees.

Overall, the results suggest that post-birth factors account for roughly 30 percent of the income gaps in Section 4, with substantial heterogeneity across family-background dimensions. This decomposition rests on the assumptions that adoptive placement is uncorrelated with adoptees' endowments and that the estimates generalize beyond the adoptee population, both of which I can assess only partially. Subject to these caveats, the results point to wealth and family structure as dimensions where post-birth factors play a relatively larger role.

⁴⁹The coefficient for migration background has even changed sign. One interpretation is that, holding all other family characteristics fixed, adopted children may benefit from being raised by parents who also have a migration background. However, this estimate is imprecisely estimated, so this result should be interpreted with caution.

Figure C1: Income Gaps Across Adoptive Family Backgrounds



Notes: this Figure shows a binscatter of predicted income rank against observed income rank for two groups of children. The blue circles represent 5,051 adopted children. The grey circles represent 3,802 own-birth children from families with at least one adopted child. The predictions come from the gradient-boosted decision tree reported in figure 2. All children were excluded from the training sample to ensure out of sample prediction. Children are sorted into ten bins based on the decile in which their predicted income falls in the national distribution of predicted income. The circles report the mean observed income rank for all children in a bin. The blue line shows the fitted regression of adopted children's income rank on their predicted income rank. (***) : $p < 0.01$

Table C1: Descriptive Statistics for International Adoptees and their Parents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income rank (y)	35,27	38,66	38,44	38,88	41,64	41,45	41,85	42,42	43,28	44,03
Predicted income rank (\hat{y})	37,93	45,36	49,06	51,99	54,32	56,71	58,9	61,2	64,13	69,74
<i>Characteristics Adoptive Parents</i>										
Parental income rank	21,68	29,13	37,62	46,71	52,82	60,89	69,43	77,71	84,82	93,42
Parental wealth rank	32,1	45,85	55,13	59,73	62,28	67,25	67,65	70,75	72,81	80,69
Highest education parents	11,26	11,81	12,73	13,25	13,92	14,59	14,94	15,56	15,68	16,53
Father suspected of crime	0,09	0,05	0,03	0,04	0,05	0,02	0,02	0,02	0,02	0,02
Health expenditures parents	4562	4254	3301	2846	3066	3269	3024	2730	2612	2497
Extended family income rank	37,79	44,55	47,56	49,01	53,54	55,51	59,37	60,27	63,74	70,45
N	505	505	505	505	505	505	505	505	505	506

Notes: each column shows descriptive statistics for a group of international adoptees from the same predicted income bin. The predicted income bins are constructed by predicting the income ranks of all adoptees using the model with all explanatory variables (as in Figure 2), ranking them from low to high, and sorting them into ten equally sized bins according to their position in the predicted income distribution of all adopted children. All cells are averages.

Table C2: The Effect of Family Background on Income: Regression Results with Adoptees

	Household income rank			
	(1)	(2)	(3)	(4)
Predicted income rank	0.279 (0.035)	0.281 (0.036)	0.279 (0.037)	0.273 (0.035)
Age at migration control		x	x	x
Gender FE		x	x	x
× Country of Origin FE			x	x
× Year of Adoption FE				x
N	5,051	5,051	5,051	5,051

Notes: each column shows results from separate regressions of adopted children’s household income rank on their predicted income rank based on their family background variables. The age-at-migration control variable is measured in months. The predicted values for income are based on gradient-boosted decision trees reported in Figure 2. All families with adopted children were excluded from the training data. The fixed effects are fully interacted. Standard errors (in parentheses) are always clustered on the country of origin level.

Table C3: Regression Results with Adoptees: Heterogeneity by Origin Country

	Household income rank					
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted income rank	0.216 (0.071)	0.365 (0.096)	0.301 (0.125)	0.274 (0.143)	0.488 (0.195)	0.312 (0.192)
Origin country	Sri Lanka	Indonesia	South Korea	Colombia	India	Brazil
N	2030	958	625	530	293	252

Notes: each column shows results from separate regressions of adopted children’s household income rank on their predicted income rank based on their family background variables. Each column shows results for adopted children from different origin countries. The predicted values for income are based on gradient-boosted decision trees reported in Figure 2. All families with adopted children were excluded from the training data. Standard errors are reported in parentheses. (***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$)

Table C4: The Relationship Between all Family Characteristics and Child Income

	Household income rank		
Parental income	0.063*** (0.001)	0.055*** (0.006)	0.016*** (0.005)
Parental wealth	0.030*** (0.001)	0.027*** (0.006)	0.012** (0.005)
Extended family outcomes	0.026*** (0.001)	0.028*** (0.005)	0.003 (0.005)
Family structure	0.021*** (0.000)	0.021*** (0.005)	0.013*** (0.005)
Parental occupation	0.013*** (0.001)	0.006 (0.005)	0.005 (0.005)
Migration background	0.010*** (0.000)	0.006 (0.007)	-0.012 (0.008)
Parental health expenditures	0.008*** (0.000)	0.009* (0.005)	0.001 (0.004)
Parental education	0.007*** (0.000)	0.002 (0.006)	0.000 (0.005)
Parental crime	0.006*** (0.000)	0.014* (0.007)	0.009 (0.006)
N	341,107	3,822	5,051
R^2	0.172	0.087	0.012
Sample	Own-birth children in test sample	Own-birth children in families with adoptees	Adopted children

Notes: columns 1 and 2 present the results of separate regressions of household income rank on 9 indices, each reflecting a different family background dimension. The sample in column 1 corresponds to all children from the test sample, as in Figure 2. The sample in columns 2 and 3 correspond to the sample of adoptees (column 3) and own-birth children in families with adoptees (column 2), as in Figure C1. The construction of the indices is explained in the text above. Standard errors are reported in parentheses. (***) : $p < 0.01$, (**): $p < 0.05$, (*) : $p < 0.1$)

Appendix D: Comparison with Cholli et al. (2024)

This appendix compares the method to estimate neighborhood intergenerational mobility in this paper to the related approach of Cholli et al. (2024). Cholli et al. (2024) use two approaches: one that controls for selection on observable family characteristics, and an extension that controls for selection on unobservables via a control function. Since this paper does not address selection on unobservables, I focus the comparison on their first approach. Cholli et al. (2024) also distinguish between family characteristics and *social* characteristics (the average characteristics of other families in the same neighborhood). For simplicity, I exclude the social characteristics from the comparison; their effects are absorbed into the neighborhood fixed effects below.

The approach in Cholli et al. (2024): Consider family background characteristics $\mathbf{X}_f = (Y_f, \mathbf{Z}_f)$, where Y_f is the parental income rank and \mathbf{Z}_f are all other observed family characteristics. Cholli et al. (2024) generalize the workhorse intergenerational mobility model by first estimating the following linear equation for the full population:

$$Y_{sf} = \eta_n + \kappa_n Y_f + \eta \mathbf{Z}_f + \kappa \mathbf{Z}_f Y_f + e_{sf} \quad (8)$$

Their measure of absolute upward mobility that corrects for sorting, which they call the location effect, is defined as follows: $\text{AUM}_{\text{location}}(n) = \hat{\eta}_n + \hat{\kappa}_n \cdot 25$.

The approach in this paper: I first estimate $\hat{Y}(\mathbf{X}_f) = E[Y_{sf} | \mathbf{X}_f]$ nonparametrically at the population level, and then run the following regression for each neighborhood separately:

$$Y_{sf} = \delta_n + \gamma_n \hat{Y}(\mathbf{X}_f) + \nu_{sf}. \quad (9)$$

I then define $\hat{Y}_{25} = \{Y : P(\hat{Y}(\mathbf{X}_f) \leq Y) = 0.25\}$ as the 25th percentile of the expected income distribution, and estimate $\text{MAUM}(n) = \hat{\delta}_n + \hat{\gamma}_n \hat{Y}_{25}$.

How the approaches differ conceptually. Equation 8 shows that Cholli et al. (2024) treat parental income Y_f as a fundamentally different family characteristic than all other family characteristics \mathbf{Z}_f . In this paper, parental income is one element of a multidimensional family-background vector \mathbf{X}_f and is treated symmetrically with the other characteristics. Four conceptual differences between the two approaches follow from this distinction.

1. *Neighborhood-specific effects.* Except for parental income, none of the family background characteristics have a neighborhood-specific slope in Equation 8. This assumes that only the relationship between parental income and child income can differ across neighborhoods. Equation 9 allows all family background characteristics \mathbf{X}_f to affect Y_{sf} in a neighborhood-specific way, with the restriction that their effects are scaled identically in each neighborhood. For example, a neighborhood with $\gamma_n = 1.2$ is one where every family characteristic matters 20 percent more than the national average. Equation 7 further relaxes this equal scaling assumption by allowing each of nine family-background dimensions to have its own neighborhood-specific slope.

2. *Non-linearities and interactions.* Equation 8 restricts \mathbf{Z}_f to enter the conditional expectation function linearly and allows for interactions only between \mathbf{Z}_f and Y_f , not among the different family characteristics in \mathbf{Z}_f . Equation 9 relaxes this restriction by estimating $\hat{Y}(\mathbf{X}_f)$ nonparametrically, accommodating arbitrary non-linearities and interactions in the population-level conditional expectation function.
3. *Evaluation point.* $\text{AUM}_{\text{Location}}$ is evaluated at the 25th percentile of the parental income distribution.⁵⁰ MAUM is evaluated at the 25th percentile of the expected income distribution. The children at the 25th percentile of either distribution have quite similar expected income (Figure 2). However, the two approaches diverge in the tails: children at $\hat{Y}(\mathbf{X}_f) = \hat{Y}_5$ have substantially lower expected income than children at $Y_f = 5$, because the multidimensional measure includes families that are also disadvantaged on dimensions other than income. When the goal is to identify how neighborhoods affect the income prospects for the most disadvantaged children—regardless of which family dimension drives their disadvantage—evaluating at percentiles of the expected income distribution may be more informative.
4. *Relative intergenerational mobility.* Lastly, both papers reports cross-neighborhood variation in relative intergenerational mobility. The standard measure of relative intergenerational mobility is the relationship between parent and child income (β_n in Equation 5). A high β_n implies that, in neighborhood n , income gaps between children from high- and low-income families are relatively large. The parameter γ_n from Equation 9 has a similar interpretation, but uses all family information rather than parental income alone to characterize a family as advantaged (high expected income) or disadvantaged (low expected income).

The intergenerational mobility parameter κ_n in Equation 8 has a different interpretation: it measures the *partial* association between parental and child income in neighborhood n , controlling for \mathbf{Z}_f . Conceptually, κ_n measures income gaps between children whose parents are similar on all observed dimensions *except* parental income.

Cross-neighborhood variation in κ_n and γ_n therefore captures different objects. Variation in κ_n captures whether the importance of specifically parental income varies across neighborhoods. If \mathbf{Z}_f includes all confounders, this would identify cross-neighborhood heterogeneity in the causal effect of parental income. Variation in γ_n instead captures whether the importance of family background more broadly varies across neighborhoods.

How the results differ empirically. Both papers report signal standard deviations of neighborhood-level mobility estimates before and after incorporating family characteristics beyond parental income, which permits a direct comparison.⁵¹ Cholli et al. (2024) report

⁵⁰Because \mathbf{Z}_f is not interacted with the neighborhood, $\text{AUM}_{\text{Location}}$ is assumed to be equal for children from all low-income parents, regardless of their other characteristics \mathbf{Z}_f .

⁵¹See Table A6 and Tables E.4 and E.5 in Cholli et al. (2024). To avoid confusion, the main decomposition in their Table 2 instead decomposes the variance of *individual-level* expected income and attributes 89 percent or more of this variance to observable family characteristics as opposed to residual neighborhood effects—which is one of their main results. This decomposition is not directly informative about how much *neighborhood-level* variation in expected income reduces once broad family information is controlled for.

reductions in the signal standard deviation of absolute upward mobility of 25 percent at the municipality level and 29 percent at the parish level under selection on observables.⁵² In this paper, the corresponding reductions are larger: the signal standard deviation of absolute upward mobility falls by 53 percent across municipalities and 55 percent across neighborhoods. This larger reduction could be explained by the broader set of family characteristics used in this paper, differences in the construction of the absolute mobility measures discussed above, or cross-country differences in the extent of sorting across neighborhoods.

⁵²For relative mobility, the corresponding parish-level reduction is from 0.084 to 0.076 (10 percent). In this paper, it falls from 0.150 to 0.071 (53 percent).