Who Gets Ahead? Measuring Income Gaps across Family Backgrounds

Sander de Vries*

December 16, 2025

Abstract

Measuring income gaps across family backgrounds is crucial for evaluating inequality and informing policy. I develop a multidimensional approach and apply it to Dutch administrative data with exceptionally detailed information on parents and extended family. This approach reveals much larger income gaps than standard intergenerational mobility measures, particularly among the most disadvantaged children. A new decomposition shows that income gaps are largest along income and wealth of parents and extended family. All family characteristics jointly can explain over half of the variation in intergenerational mobility across neighborhoods. These results matter for policymakers that seek to target disadvantaged families or 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

Concerns about inequality often focus on the role of family background in shaping children's economic success. Low income prospects for children from disadvantaged families are widely seen as a failure of equal opportunity and as a justification for policy intervention (Roemer and Trannoy (2016), Alesina et al. (2018)). Yet, despite its importance, measuring the size and nature of such income gaps remains challenging.

Recent administrative data have allowed researchers to map in detail how parental and child incomes are related (e.g., Chetty et al. (2014), Deutscher and Mazumder (2020), Kenedi and Sirugue (2023)). These studies reveal which children have the lowest and highest expected incomes and their chances of moving up or down the income distribution. Yet 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. Analyses based on parental income alone miss these disparities. While measures such as sibling correlations or inequality of opportunity estimates aim to capture such broader influences, they typically summarize it in a single statistic and therefore conceal how large the income gaps are between specific types of families.[1] As a result, two important questions remain unresolved: how large are income gaps across broader family background characteristics, and what distinguishes the families where children have systematically low or high incomes?

This paper sheds light on these questions by addressing two main challenges. The first is that a rich characterization of family background requires linked child–parent data with information across many dimensions. I draw on Dutch administrative data that link the long-run incomes of 1.7 million children to detailed information on their parents and extended family, including income, wealth, occupation, education, crime, healthcare, migration background, and family structure. This makes it possible to examine disparities between families

[1]Deutscher and Mazumder (2023) distinguish between 'global' measures, which summarize the overall role of family background, and 'local' measures, which provide distributional detail. They also separate 'narrow' measures, based on a single parental characteristic, from 'broad' measures that capture the broader family background. They show that no existing study offers a broad but local analysis.

in much greater detail than previous work.

The second challenge is to translate this high-dimensional information into interpretable measures that reveal fine-grained distributional patterns. 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. However, 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 parts of the income distribution. In addition, I adapt a new method based on Shapley values to describe how different family background characteristics contribute to *each* child's expected income rank (Lundberg et al. (2020)). Unlike traditional variance decomposition methods, this method reveals family background characteristics that are important for some children even if their aggregate contribution is modest.

As a starting point, I examine the relationship between child and parental income, following the standard intergenerational mobility approach. I estimate a rank-rank correlation for the Netherlands of 0.32 (R^2 = 10.5 percent), which is high compared to other OECD countries. This shows that parental income alone already predicts substantial income gaps.

I then show that focusing on parental income alone severely understates income disparities across family backgrounds, especially for the most disadvantaged children. Incorporating all family information raises explanatory power by 58 percent, with these family characteristics jointly accounting for 16.6 percent of income variation. The difference between the two approaches is greatest at the bottom of the distribution. For instance, the 0.5 percent of children with the lowest expected incomes based on parental income alone have an average observed rank of 31 (out of 100). With all family information, this falls to 18. Their chance of reaching the top income quintile is below two percent. These findings show that a multidimensional approach can uncover large income gaps that remain hidden when the analysis relies on a single variable or summarizes disparities into a single index.

Using the Shapley value decomposition, I show how much the family characteristics contribute to expected income. Nine of the ten variables that explain most income variation relate to parental and extended family income and wealth, underscoring the central role of economic resources and extended kinship in quantifying income gaps across family backgrounds. The granularity of the decomposition also makes it possible to identify family characteristics that strongly affect expected income but explain relatively little income variation. Parental absence, for instance, has low variation and thus contributes little to aggregate income inequality, yet the decomposition reveals that it substantially lowers expected income for affected children. The low expected income of the most disadvantaged children arises from the cumulative negative contributions of many variables: their parents are often young, separated, and have low income and wealth, limited education, high health expenditures, and criminal records, with similar disadvantages among aunts and uncles.

Lastly, I show that the income prediction model is useful for estimating income gaps across neighborhoods and among international adoptees, where samples are too small to estimate similarly rich models non-parametrically. Expected income of children from low-income families varies widely across neighborhoods, but this variation falls by over half once I compare children who are similarly disadvantaged based on all family information. This shows that a sizable share of differences in intergenerational income mobility across neighborhoods reflects differences in underlying family characteristics. The analysis with adoptees suggests that pre-birth factors play an important role in driving income gaps across family backgrounds. Being raised from infancy in an advantaged family increases the income of adoptees, but by considerably less than what is predicted for own birth children.

This paper offers three contributions. First, I provide the most detailed measurement of income gaps across observable family background characteristics to date. The population-wide data allows me to examine disparities even among very small and (dis)advantaged groups of children, while the richness of the family information makes it possible to evaluate the relative importance of many family background dimensions that were previously analyzed in isolation. Prior work linking multiple family background characteristics to child income often collapses the income disparities into a single statistic, such as intergenerational mobility coefficients (Vosters and Nybom (2017), Vosters (2018), Adermon et al. (2021), Eshaghnia et al. (2022), Chang et al. (2025)), inequality of opportunity estimates (Brunori et al. (2023), Adermon et al. (2025)), or explanatory power measures (Blundell and Risa (2019), Althoff et al. (2025)).[2] These summary statistics are very useful for making comparisons across regions or over time, but a drawback is that they mask distributional patterns. Providing distributional insights matters because fairness concerns and policy choices also depend on where in the distribution of family backgrounds the income gaps are largest.[3]

Second, I use the novel decomposition by Lundberg et al. (2020) to quantify how much the family background characteristics contribute to each child's expected income. Earlier work has measured variable contributions to summary statistics such as explanatory power or inequality of opportunity estimates (Mendolia and Siminski (2017), Blundell and Risa (2019), Salas-Rojo and Rodríguez (2022), Brunori et al. (2023), Althoff et al. (2025)). These approaches measure average contributions. Decomposing predictions at the child level also shows for how many children, and which children, these variables have a meaningful effect. Although this method is widely used in machine learning and is well suited for decomposing inequality, it has not, to my knowledge, been used to study individual income differences.

Third, I contribute to the literature on spatial differences in intergenerational mobility.[4] Chetty et al. (Forthcoming) use income predictions based on the national relationship between parental and child income to estimate neighborhood upward mobility. I extend their approach by using income predictions based on all family information. This considerably reduces the variation in expected income across neighborhoods, which implies that a sizeable share of differences in upward mobility across neighborhoods can be explained by differences in these previously unobserved family characteristics.[5] The adjusted estimates remain

[2]Brunori et al. (2023) and Brunori et al. (2024) also provide distributional results, but rely on much smaller survey data with fewer family background variables.

[3]For instance, even when summary statistics are equal, a small group of highly disadvantaged children imposes different fairness and policy considerations than a small group of highly advantaged children.

[4]See Chyn and Katz (2021) and Mogstad and Torsvik (2023) for broad overviews of this large literature.

[5]Cholli et al. (2024) reach a similar conclusion with a different approach for Denmark. Using a control-function approach that includes detailed family and community characteristics, they show that sorting explains much of the variation in intergenerational mobility across neighborhoods.

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

This paper proceeds as follows. Sections 2 and 3 present the methodology and the data. Section 4 measures income gaps across family backgrounds and discusses the most important predictors. Sections 5 and 6 use the comprehensive prediction model to measure income gaps across neighborhoods and among international adoptees, respectively. Section 7 concludes.

2 Methodology

This section illustrates how multiple background characteristics can be used to quantify income gaps across family backgrounds. I begin with a simple and widely used summary statistic, and then introduce two measures that provide more fine-grained insights. Finally, I show how recent advances in machine learning can be used to assess the relative importance of different family background variables for specific individuals.

Global measure. Let Y_{sf} be the income rank of a child s in a family f. Moreover, let $\mathbf{X}_f = (X_{f1}, ..., X_{fK}) \subset \mathcal{X}$ be the set of all observable family background characteristics. Consider the following conditional expectation function decomposition of Y_{sf}:

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

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

The primary objective of this paper is to measure the importance of observable family background characteristics for children's 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 non-parametric R^2 of the observables model:

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

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).[6]

Another 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.[7] 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

[6]This 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 includes factors that differ between siblings, such as birth order effects or life-cycle variations in parental income over time.

[7]A 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.

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.

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, non-parametric machine learning methods outperform linear regression models. Accordingly, I employ gradient-boosted decision trees to generate these predictions (Friedman (2001)).[8] 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. After estimating expected income for each child, a next question is how much each of the family background characteristics contribute to

[8] Single 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)).

differences in expected income between children.

Decomposing family characteristics' contributions to expected income is not trivial. For example, based 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 full model may contribute substantially in smaller subsets.

Instead, 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.[9] 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)).[10] 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 $S \subseteq \mathbf{X}_f \backslash \{X_{fk}\}$ by the change in expected income induced by adding this variable:

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

[9]These 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.

[10]Their 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.

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 $h(X_{fk}, \mathbf{X}_f \backslash \{X_{fk}\}) = g(\mathbf{X}_f) - E[Y_{sf}|\mathbf{X}_f \backslash \{X_{fk}\}]$.

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 of the covariates:

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

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.

Shapley values are an attractive way to assign 'contributions' to variables because, for each child, the Shapley values of all variables sum to this child's expected income:

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

Importantly, 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.

Because this approach provides Shapley values for each child, I can show in detail how much and for how many children each of the variables contributes to their expected income. Prior approaches typically compute Shapley values for aggregates, such as a model's explanatory power or inequality (of opportunity) indices (Shorrocks, 2013). These decompositions are informative about which variables contribute most to the aggregate measure but offer limited insight into which factors matter for particular individuals or subgroups. This is especially relevant in the analysis of inequality, where we may wish to understand what drives the high or low expected income of smaller groups of individuals at the tails.

3 Data

Core analysis sample. I use administrative data from Statistics Netherlands covering the full Dutch population.[11] The main sample consists of all children born in the Netherlands between 1980 and 1989, excluding 3.4 percent with missing income data, resulting in 1,703,038 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 incomes are observed between 2003 and 2023 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.[12] 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 income from age 30 onward.[13] This approach reduces measurement error from transitory

[11] Access is granted through a secure remote facility under a confidentiality agreement.

[12] Some 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.

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

income shocks (Mazumder, 2005) and life-cycle bias (Haider and Solon (2006), Nybom and Stuhler (2017)). I observe income up to age 43 for the oldest cohort (born in 1980) and up to age 34 for the youngest cohort (born in 1989). On average, children have nine 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 2003 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 12 income observations and mothers 14. 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 eight 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 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

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, household income rank, wealth rank, total 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 B.
*: 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).

years. Nonetheless, many parental characteristics are highly persistent over the life cycle, making them a reasonable proxy for the family environment at earlier ages.[14]

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 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.

[14]This is supported by Eshaghnia et al. (Forthcoming), 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.

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 representing half a percentile and 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 income ranks but increases steeply at the tails of the parental income distribution.[15] An OLS regression yields a slope coefficient of 0.32, indicating that a one-rank increase in parental income corresponds to a 0.32-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 relatively strong persistence. Its rank-rank correlation is higher than in Sweden, Denmark, Australia, Norway, and Canada (0.20–0.24), similar to France, Germany, Italy, and the UK (≈ 0.30), and below the United States (0.36). This is striking 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 measurement error and life-cycle bias.

[15] As 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.

[16] This 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 our estimates differ.

Figure 1: Intergenerational Income Mobility in the Netherlands

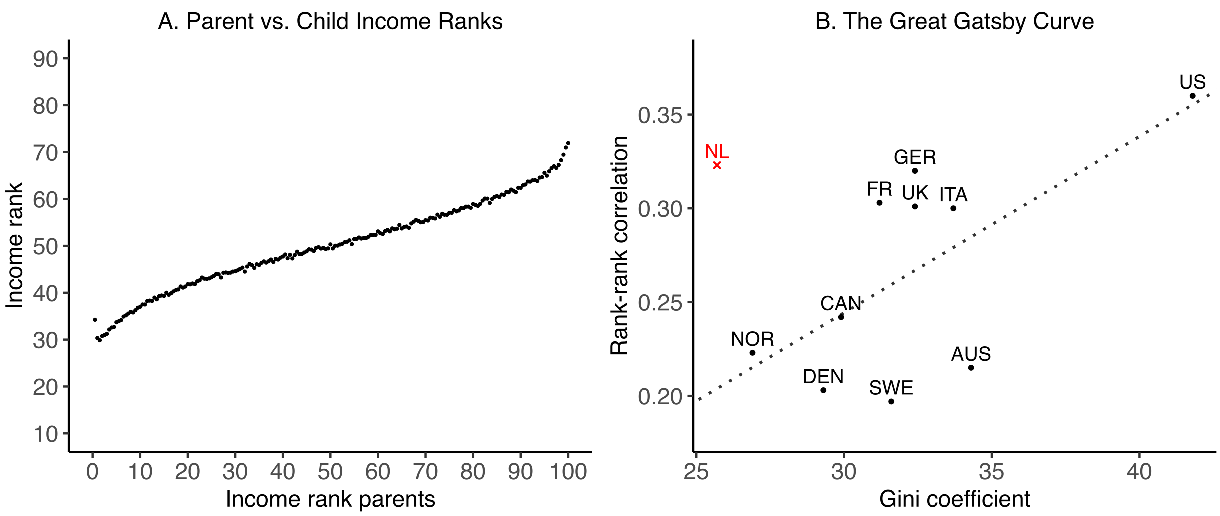

Notes: panel A presents a nonparametric scatter plot of mean income ranks versus parental income rank. The sample consists of all $N = 1,702,355$ children from the core analysis sample (Table A1) for whom parental income is not missing (99.1%). The X-axis reports the parent income rank sorted into 200 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.

4.2 Including Detailed Parental Information

Explanatory power. To quantify the increase in family-driven inequality when adding the broader family background information, I first compare the explanatory power of a model using only parental income with that of a model incorporating all explanatory variables. Both models are trained and evaluated on the same training and test data. For the income-only model, I non-parametrically predict a child's income rank in the test data by the mean income rank of all children in the training data with the same parental income rank and year of birth. Like the linear regression in the previous section, this model achieves an explanatory power of 10.5 percent. The predictions using all explanatory variables are generated by a tuned gradient-boosted decision tree, as described in Section 2. This model includes all explanatory variables from Table 1 and children's year of birth.

Adding all information about the parents reveals substantially stronger intergenerational dependence. The comprehensive model achieves an explanatory power of 16.6 percent, marking a 58 percent increase compared to the income-only model (Figure 2).

To put this into perspective, an increase in the rank-rank correlation from 0.32 to 0.41 would result in the same increase in R^2.[17] This is considerable, considering the difference in rank-rank correlation between Denmark (high mobility) and the US (low mobility) is about 0.16 (Helsø (2021), Davis and Mazumder (2024)). Moreover, the increase in R^2 far exceeds the gain achieved from reducing attenuation bias in an income rank-rank regression, a source of measurement error that has received considerable attention in the literature (Mazumder (2005), Nybom and Stuhler (2017)).[18] When the goal is to quantify income disparities between families, adding more information about parents is thus more valuable than constructing a more accurate proxy of lifetime income.

Distribution of expected income. I next explore where in the income distribution the

[17] I use here that in a rank-rank regression, $R^2 = \beta^2$ (i.e. $0.408^2 - 0.324^2 = 0.166 - 0.105 = 0.061$).

[18] Table B2 columns 1 and 9 shows that using 9 years of income data versus one year of income data in a rank-rank regression increases the R^2 from 8.2% to 10.0%.

Figure 2: Predicting Child Income with Detailed Parental Information

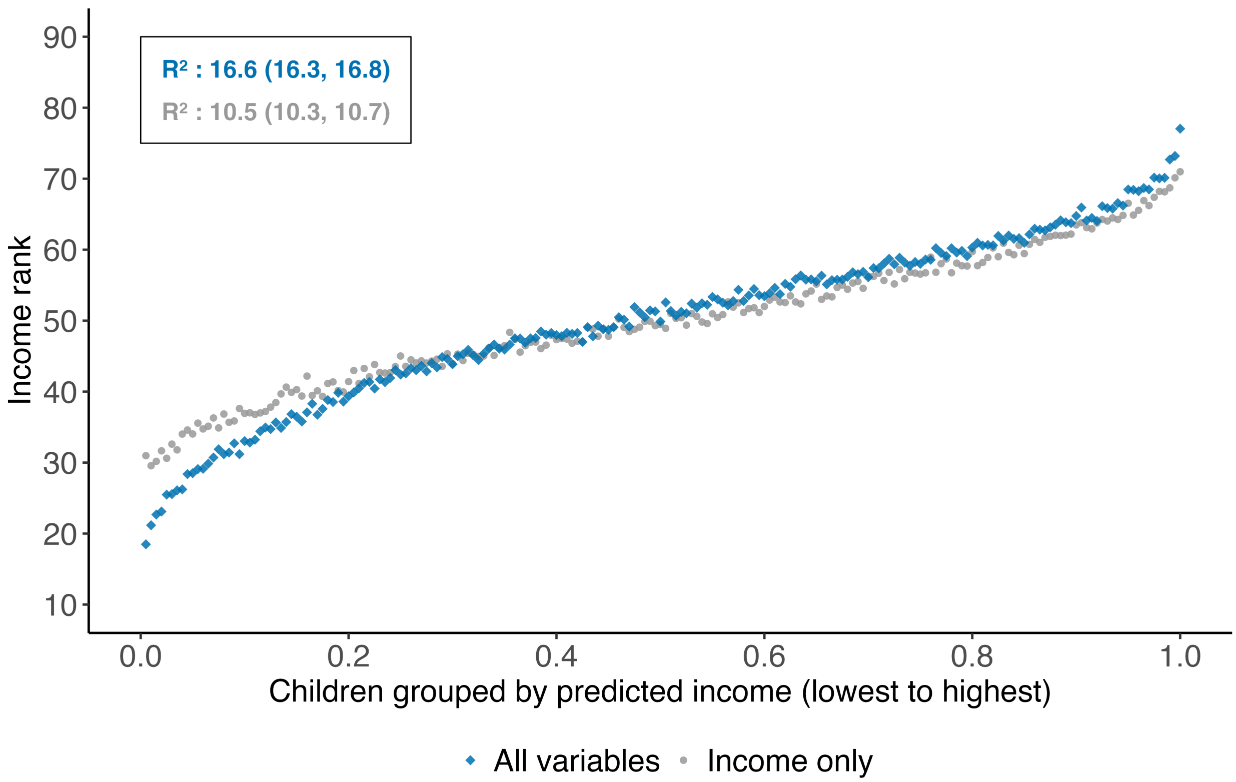

Notes: this Figure presents binscatter plots of income ranks for 340,608 children in the test data, who are sorted into bins based on their predicted income rank according to two models. Both models are trained to predict children's income ranks using the same training sample of 1,362,430 children but 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 graphs are 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.

difference between the two approaches is greatest. Figure 2 provides a detailed binscatter plot of children's mean income ranks, sorted from lowest to highest predicted income. The X-axis divides the test dataset into 200 bins, each containing approximately 1,700 children, based on their predicted income ranks within their cohort. The Y-axis reports the average observed income rank for each bin. The blue dots represent children grouped by predicted income using parental income alone, while the grey diamonds reflect groupings based on predictions from the comprehensive model.

The comprehensive model identifies considerably greater income disparities by family background, particularly for the most disadvantaged children. For instance, in the income-only model, the 0.5 percent of children with the lowest expected income have an average income rank of 31. With the comprehensive model, this drops to 18. Similarly, for the top 0.5 percent, the income-only model estimates an average rank of 70, while incorporating additional family background information raises this to 78. To the best of my knowledge, no other study has identified children with similarly low or high expected income ranks based solely on family background information. This demonstrates that a multidimensional approach can uncover large and previously unseen income gaps across family backgrounds.

Transition probabilities. Figure 2 reports only children's average income rank, yet there is substantial variation around this average. 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 any income quintile should be 20 percent for all individuals. Instead, there are large differences. For instance, among the 0.5 percent children with the lowest expected income, less than two percent are in the top income quintile, whereas 69 percent are in the bottom income quintile.

To place these 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 nearly four times more likely to be in the bottom quintile than the top (35.8 vs. 9.3 percent).[19] Incorporating the full set of family background variables increases this ratio to almost seven (41.8 vs. 6.1 percent). These results indicate that for a substantial share of Dutch children, the chances of 'moving up' are strikingly low, and much lower than implied by parental income alone.

4.3 What Characterizes Family (Dis)Advantage?

To understand which of the included family characteristics are most strongly associated with child income, I use the Shapley value decomposition described in section 2. Figure 4 presents a detailed graph illustrating the variable importance of the 30 most predictive variables, calculated using Shapley values.[20] The boxplots report the distribution of Shapley values for a randomly drawn subset of 10,000 children from the test data. As they are randomly drawn from the full population, these estimates are representative of the full population.

To illustrate their interpretation, consider the most important predictor of child income: parental income. The 2.5th percentile of the corresponding boxplot is -6.5, indicating that for 2.5 percent of the children, the Shapley value for parental income is below 6.5. As discussed in Section 2, a Shapley value of -6.5 means that, for this specific child, including parental income to a set of other predictors on average reduces expected income by 6.5 ranks, where the average is taken over all possible subsets of predictors. The 75th percentile is 3.3, meaning that for 25 percent of the children, adding parental income increases expected income by more than 3.3 ranks. Relative to the other variables, the distribution of Shapley values for parental income is wide. This means that parental income information considerably changes

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

[20] All variables contribute to the predictions, but the effects of the remaining variables are small and are therefore omitted for conciseness.

Figure 3: Child Income Distributions Across Family Backgrounds

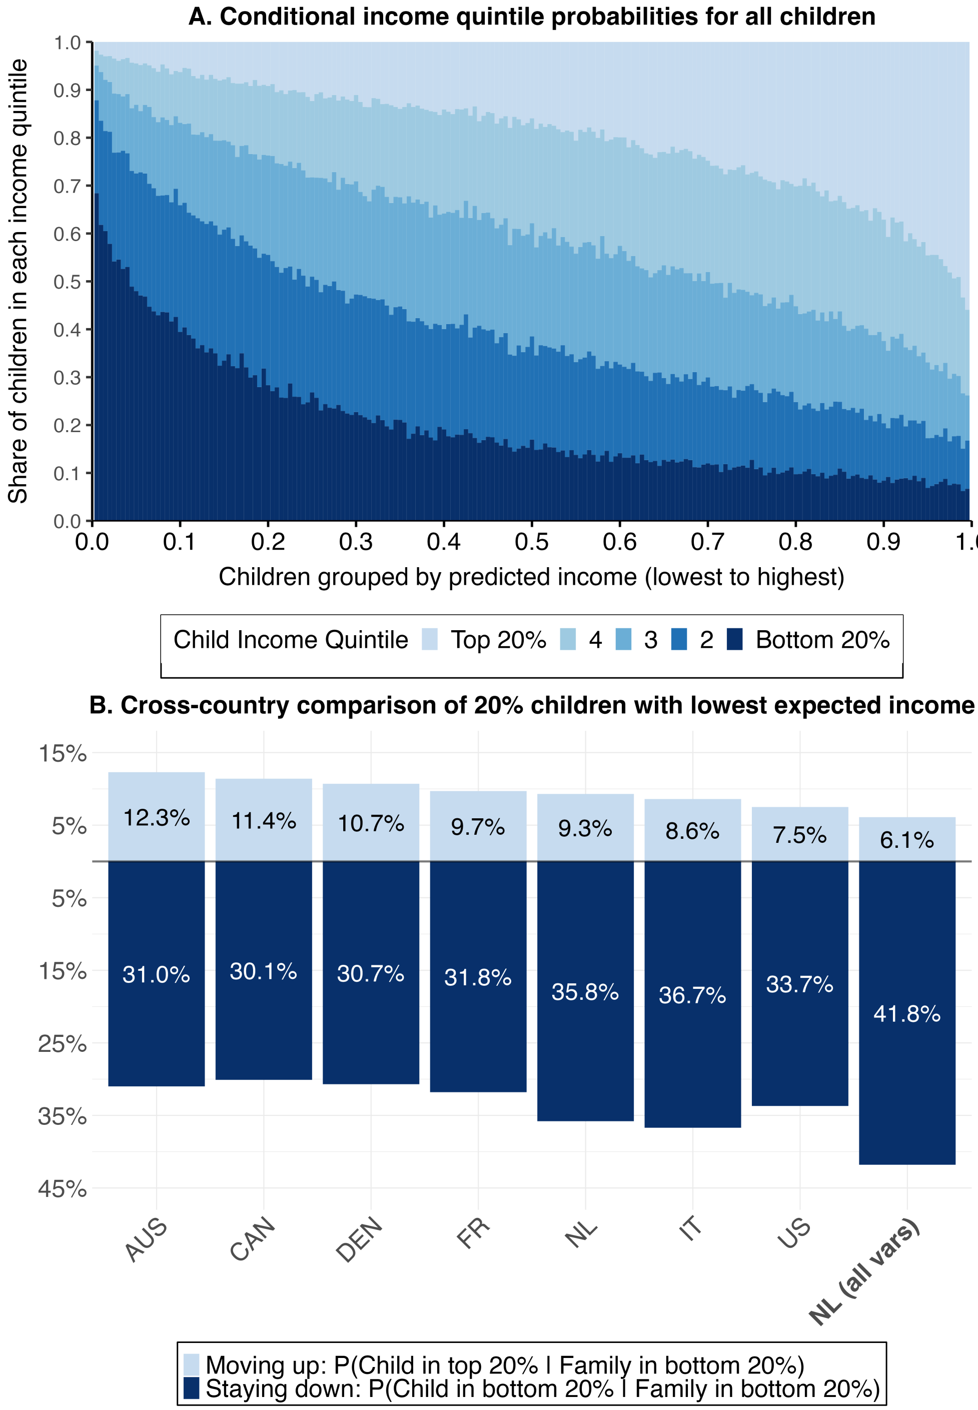

Notes: figure A shows the share of children in each income quintile. The sample contains all 340,608 children from the test data, and they are grouped into 200 equally sized bins according to their predicted income rank from the comprehensive model (as in Figure 2). Figure B shows the share of children 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).

the expected income of many children, both upward (for children with high parental income) and downward (for children with low parental income).

Figure 4 orders the variables by their average absolute Shapley value. This ranking defines the most important predictors as those with the largest contributions, both positive and negative, to the predictions. Nine of the ten most important predictors are all related to parental or extended family income and wealth. This underscores the importance of economic resources and extended family ties in capturing family-driven income inequality.

While other variables contribute less on average, they can make sizeable changes to the expected income of smaller subsets of children. To illustrate this, consider the Shapley values for mother presence. The box is centered around zero, indicating small contributions to expected income for most children. This is expected, since 96 percent of children live with their mother at age 15. However, the left whisker shows that for 2.5 percent of children, whose mothers are absent, expected income decreases by at least 2.4 ranks. This shows that although mother absence contributes little in the aggregate, it substantially reduces the expected income of a small group. Such effects would be overlooked by methods which focus only on contributions to aggregate explanatory power.

Another advantage of individual-level Shapley values is that they can provide insights for specific groups of children. The triangles and diamonds in Figure 4 report the average Shapley value for each variable for the 0.5 percent of children with the lowest and highest expected income, respectively. Consider the most disadvantaged children. Their expected income rank is close to 18, which implies that the sum of their Shapley values must be around -32 (Equation 4). The average Shapley value for parental income is -7, which is sizeable but still far away from -32. This implies that many variables contribute jointly to their low expected income. This is illustrated in Figure 4 by the large set of family characteristics with negative average Shapley values, including those with limited aggregate importance such as living in a single parent household or having a father suspected of a crime. These moderate contributions across many variables accumulate into the very low expected income

Figure 4: The 30 Most Predictive Family Characteristics of Child Income

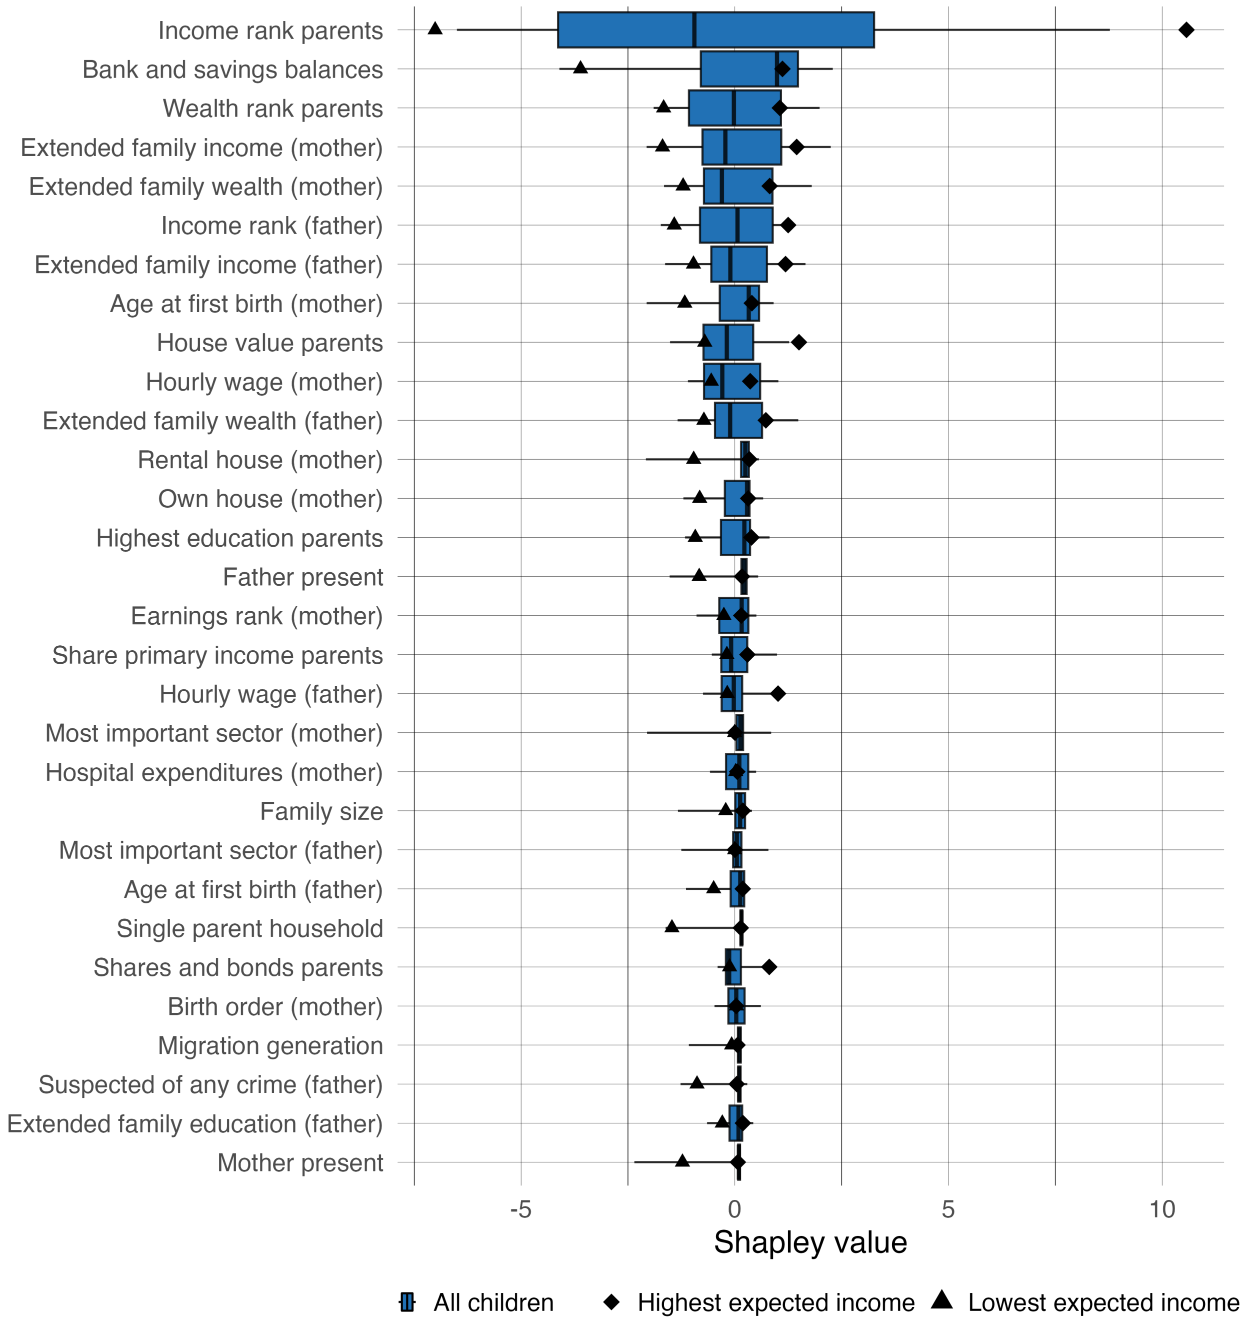

Notes: each row displays a boxplot representing the distribution of Shapley values for a given variable. These Shapley values are computed using the algorithm of Lundberg et al. (2020) for each variable and each child using a randomly drawn sample of 10,000 children from the test dataset. The variables 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. The triangles (diamonds) report the average Shapley value for each of those variables for the 0.5 percent of children ($n = 1703$) from the test data with the lowest (highest) predicted incomes.

ranks observed for these children.

To further illustrate these cumulative (dis)advantages, Table A2 provides descriptive statistics of children with different expected incomes. This table shows that the family characteristics of the children at the bottom of the expected income distribution are unfavorable in nearly every measurable dimension. They have parents with low income and wealth and who are often young, separated, minimally educated, suspected of crimes, have high health expenditures, and their aunts and uncles also have low income and wealth.

4.4 Additional Results

Sibling correlation. A commonly used alternative method to quantify the importance of family background is the sibling correlation, which captures the contribution of all factors shared between siblings (Solon (1999)). I estimate a sibling correlation in income of 0.308 (Table B1). As discussed in Section 2, 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 about half of this correlation (0.166/0.308). The remaining half of siblings' similarities may be explained by other shared factors, such as 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.

Income level differences. The findings above relate to inequalities in gross household income ranks. In Figure A3, I provide results using gross household income and disposable household income in absolute terms, while maintaining the same explanatory variables.[21] The comprehensive model's explanatory power for household income is 11.2 percent, which is higher than the 8.5 percent for disposable income. This shows that income redistribution through taxes and transfers somewhat diminishes the impact of family (dis)advantages.

Predicting education and violent crime. Figure A4 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 12.5 percent to 25.6 percent, an increase of 103 percent. For violent crime it rises from 3.9 percent to 10.5 percent, an increase of 169 percent. Both gains far exceed the 58 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 15.3 percent.[22] This is quite close to the explanatory power of the comprehensive model, suggesting that incorporating a broader range of information is more critical than allowing for complex interactions and non-linearities. However, a downside of OLS regression is that it does not support the computation of individual-specific Shapley values. While 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.

Robustness. Table A3 shows that explanatory power declines with smaller samples but stabilizes once at least 40 percent of the data are used. This suggests that downward bias

[21]Disposable income is the amount left after deducting taxes and social insurance payments from gross income.

[22]Coefficient estimates are available upon request.

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. Explanatory power attenuates when fewer years of income are used, but stabilizes once about five years of income are used.[23] 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 virtually identical 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 ($n = 1,093,245$, $R^2 = 17.4$). I then re-train the model on the same sample after removing all education variables for both parents and extended family ($R^2 = 17.3$). 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.

5 Measuring Income Gaps Across Neighborhoods

Previous research has shown that children from *low-income* backgrounds can have vastly different expected income based on the place where they grew up (e.g., Heidrich (2017), Deutscher and Mazumder (2020), Corak (2020), Alesina et al. (2021), Acciari et al. (2022), Kenedi and Sirugue (2023), Chetty et al. (Forthcoming)). This subsection proposes a simple approach that incorporates additional family information into the measurement of such neighborhood disparities.

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 2,828 neighborhoods in which they

[23] In 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.

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. To estimate differences in intergenerational mobility across neighborhoods, I follow Chetty et al. (Forthcoming). They estimate for each neighborhood the expected income of a child in the 25th percentile of the parental income distribution. Because there are often too few children in a neighborhood to estimate this non-parametrically, they use a univariate regression whose functional form is chosen based on estimates at the national level. That is, for a child s in family f with parental income rank Y_f in neighborhood n, they consider the following specification:

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

where $\hat{Y}(Y_f)$ non-parametrically estimates $E[Y_{sf}|Y_f]$ based on the national-level relationship between Y_{sf} and Y_f.[24] This specification summarizes the conditional expectation function in each neighborhood using just two parameters, α_n and β_n. *Absolute upward mobility* (AUM) for neighborhood n is then defined by $AUM(n) = \alpha_n + \beta_n \hat{Y}(25)$.

I generalize this approach to include additional family information. I use the same univariate regression, but replace the predictions $\hat{Y}(Y_f)$ using parental income only by predictions

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

$\hat{Y}(\mathbf{X}_f)$ based on all family information:

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

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

I then use expression 6 to compute a new measure of upward mobility. Let $\hat{Y}_{25}$ be the 25th percentile of the national expected income distribution: $\hat{Y}_{25} = \{Y : P(\hat{Y}(\mathbf{X}_f) \leq Y) = 0.25\}$. *Multidimensional absolute upward mobility* (MAUM) is then defined 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: $MAUM(n) = \delta_n + \gamma_n \hat{Y}_{25}$. This measure allows me to compare the expected income of children who are similarly disadvantaged based on all family information, but who grew up in different neighborhoods.[25]

To operationalize these measures, 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 the 80 percent in the training sample, to increase the number of observations in each neighborhood. A cross validation procedure ensures that predictions for these children remain out of sample.[26] I then estimate $MAUM(n)$ and $AUM(n)$ by running regressions 5 and 6 for each neighborhood and plugging in the resulting estimates for α_n, β_n, δ_n, and γ_n. 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 reports histograms for the estimates of upward mobility. There is substantial variation in absolute upward mobility, with estimates ranging from 30 to 60 ranks. The standard deviation is 4.7 ranks. Taken at face value, this suggests that a one standard

[25]When parental income is the only predictor (i.e., $\mathbf{X}_f = Y_f$), then $MAUM(n) = AUM(n)$. This follows because in that case, children whose parents rank 25th in the national income distribution rank 25th themselves in the expected income distribution.

[26]Table A5 shows that the R^2 of the predictions for the full sample is 16.5 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.

Figure 5: Variation in Absolute Upward Mobility Across Neighborhoods

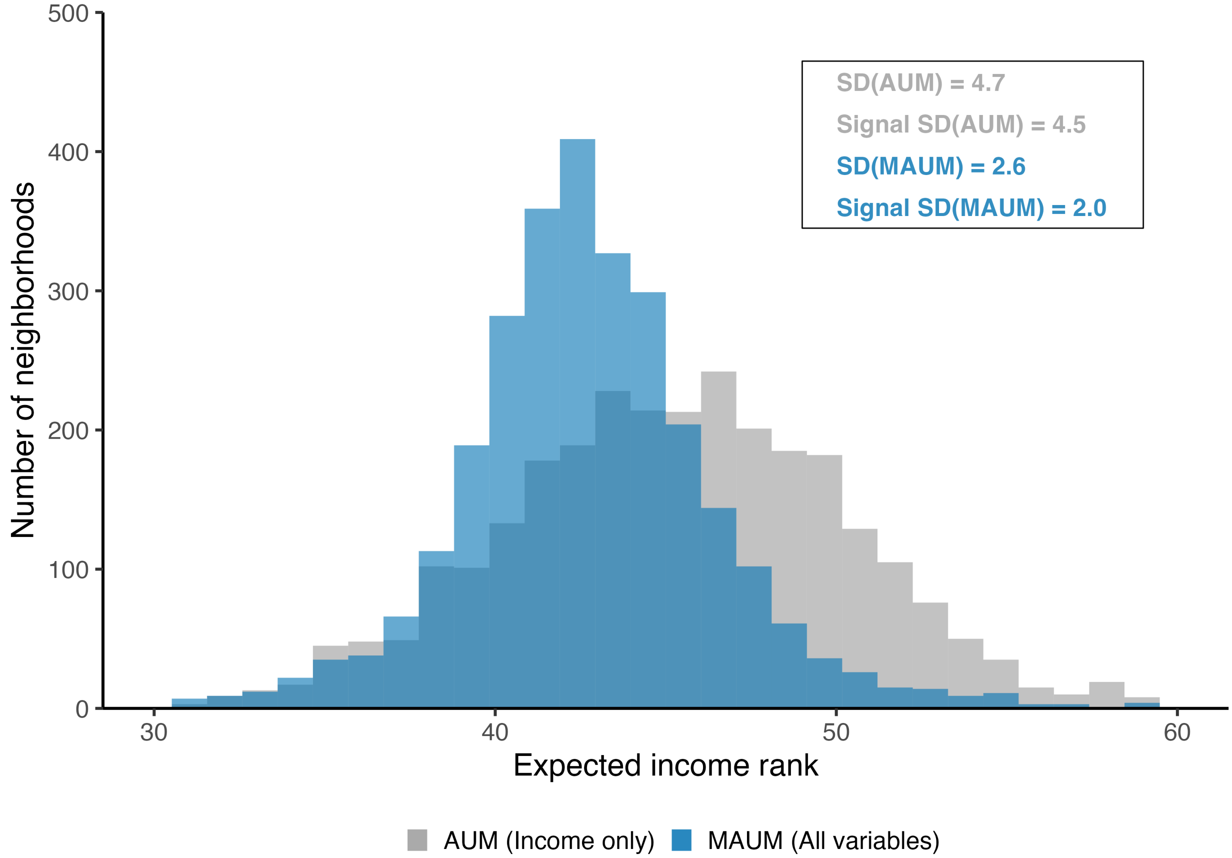

Notes: this Figure presents histograms of neighborhood-specific upward mobility estimates. The Absolute Upward Mobility (AUM) estimates are obtained by the fitted values of regression 5 after plugging in $Y_f = 25$. The Multidimensional Absolute Upward Mobility (MAUM) estimates are obtained by the fitted 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.

deviation higher upward mobility neighborhood increases expected income by 4.7 ranks. However, part of the standard deviation in the estimates reflects noise induced by sampling uncertainty, since there are on average 580 observations per neighborhood. I use the estimates' standard errors to conduct a standard signal-noise decomposition, and estimate that the signal standard deviation equals 4.5 ranks.[27] This is somewhat lower than the variation across US census tracts, where the signal standard deviation in upward mobility equals 6.2 ranks (Chetty et al. (Forthcoming)).

The multidimensional estimates show much lower dispersion.[28] Their standard deviation is 2.6 ranks, which is 45 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 42 percent. The signal noise decomposition confirms this: the signal standard deviation is 2 ranks, which is 56 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 measured with the full set of family information rather than parental income alone.

The reduced standard deviation of 2 ranks leaves a relatively modest role for neighborhoods in driving income gaps. It implies that, holding 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 11.2 ranks.[29] This shows that income gaps are considerably larger across family backgrounds than across neighborhoods.

[27] Specifically, 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.

[28] The average is also lower. This is because the average child in the 25th percentile of the expected income distribution based on all variables has a lower expected income than the average child in the 25th percentile of the parental income distribution (see Figure 2).

[29] The standard deviation of $\hat{Y}(\mathbf{X}_f)$ is 11.7 ranks. Within neighborhoods, a one rank increase in $\hat{Y}(\mathbf{X}_f)$ increases income by 0.96 ranks (Table A5). Keeping the neighborhood constant, a one standard deviation increase in expected income $\hat{Y}(\mathbf{X}_f)$ therefore raises expected income by $0.96 \times 11.7 = 11.2$ ranks.

Supplementary results and discussion. I report several alternative estimates in Table A6. 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* intergenerational mobility. The dispersion in these estimates also falls sharply 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.[30] Even at this higher level of aggregation, the dispersion still falls by 40 percent. The estimates are also stable when I estimate regressions 5 and 6 with second- or third-order polynomials to allow for nonlinearities. These results show that the decline in dispersion is robust across multiple mobility measures, a higher level of aggregation, and more flexible functional form choices.

Policymakers and families care not only about levels of upward mobility but also about the ordinal ranking of neighborhoods. Figure A5 shows that the upward mobility ranking of all neighborhoods can change substantially with the multidimensional measure. Neighborhoods who are in the same percentile with the income-only measure are on average 20 percentiles apart with the new measure.[31] As an example, 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 6 percent lowest mobility neighborhoods. The highest 10 rise above the 55th percentile. Neighborhoods which are similarly mobile based on the income-only measure can thus be ranked very differently based on the multidimensional measure.

Lastly, I show that sorting explains the larger dispersion of the income-only estimates. As shown in the main results, some low-income families may be advantaged in other dimensions. If this is true for many families in a neighborhood, then this neighborhood will exhibit high absolute upward mobility. To test for such sorting, I select all children with parental income between the 20th and 30th percentile. I then regress their predicted income $\hat{Y}(\mathbf{X}_f)$ on neighborhood absolute upward mobility $\widehat{\mathrm{AUM}}$, and include parental income as a control

[30] The 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.

[31] This is measured by the within-group standard deviation of the new ranks. Groups are neighborhoods that fall in the same percentile based on the old ranking.

variable. If low-income families in high and low upward mobility neighborhoods were similar along all other dimensions, then upward mobility would be unrelated to predicted income because predicted income is a function of only family characteristics. However, Table A7 shows a highly significant coefficient of 0.67. This means that based on children's family characteristics alone we would already expect children in 1-rank higher upward mobility neighborhoods to have 0.67 ranks higher income.

Prior studies often use (quasi)random moves to measure the share of the variation in upward mobility that reflects such sorting (Chetty et al. (2016), Chetty and Hendren (2018), Kawano et al. (2024)). For example, Chetty et al. (Forthcoming) estimate that about 40 percent of the variation across US census tracts reflects sorting. The advantage of using random moves is that it accounts for all family background differences, including unobserved ones. However, a limitation is that the number of moves is often too small to produce reliable new neighborhood specific estimates. The results above show that adjusting for a broad set of *observable* family background characteristics can already correct for large sorting effects 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. Future work could combine the multidimensional approach with quasi random moves to assess how much of the remaining variation still reflects unobserved family background differences.[32]

6 Measuring Income Gaps Among Adoptees

This last section measures income gaps among international adoptees. This provides insights into the role of post-birth factors in driving the observed disparities in Figure 2.

I consider a sample of 5,044 international adoptees born between 1980 and 1989 and who arrived in the Netherlands within six months of birth.[33] These children are not genetically

[32] This is not possible here because I do not observe addresses early enough.

[33] Although the Netherlands lacks an adoption register, Statistics Netherlands developed a reliable method

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 the background characteristics of their adoptive families. Figure 6 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 43, even though their adoptive family income, wealth, education, and other characteristics are among the highest in the population (Table A8).

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.

I do not have access to information on matching procedures from this period, restricting a comprehensive assessment of assumption 1. However, Appendix A9 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.[34]

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 ($n = 787$).

[34]The 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

Figure 6: Income Gaps Across Adoptive Family Backgrounds

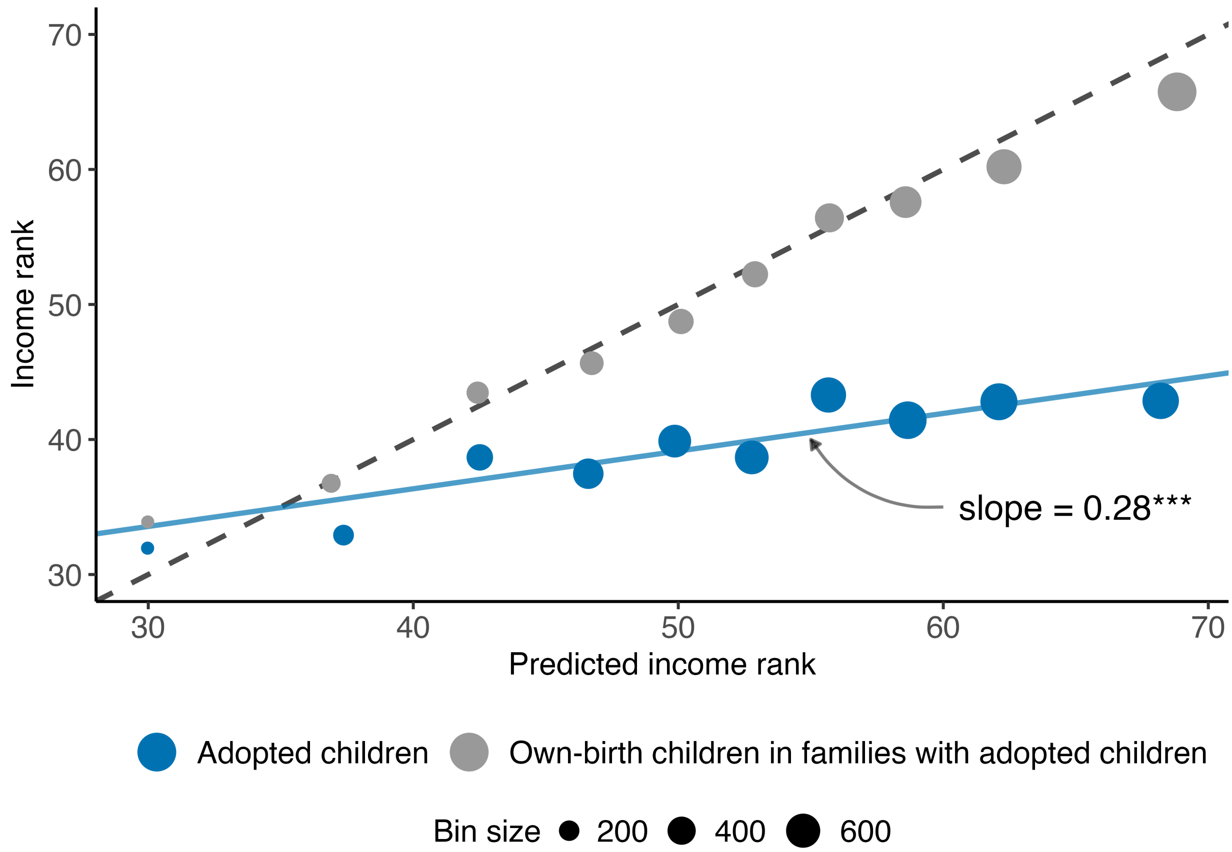

Notes: this Figure shows a binscatter of predicted income rank against observed income rank for two groups of children. The blue circles represent 5,044 adopted children. The grey circles represent 3,802 own-birth children from families with at least one adopted child. The predictions come from the gradient-boosted 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. ($^{***}: p < 0.01$)

To assess external validity, Figure 6 also reports results for 3,802 own-birth children from families with at least one adopted child. In this group, observed income aligns closely with predicted income. This indicates that intergenerational dependence in these families is similar to that in the broader population when parents and children are biologically related. I also test for heterogeneity among adoptees from the six largest origin countries in Table A10 and do not reject equality of the coefficients (p = 0.86). This suggests that the much weaker link between adoptees' income and the characteristics of their adoptive families is a robust phenomenon across migration backgrounds. Nevertheless, extrapolation of the results is difficult when adoptive parents raise adopted children differently or when adopted children's initial endowments such as infant health differ.[35] The importance of these issues remains an open question.

Lastly, in Appendix C, I examine which adoptive family characteristics are most strongly associated with adoptees' income. This is challenging because estimating a new model with this large number of predictors and relatively small sample of adoptees is infeasible.[36] I therefore use Shapley values to summarize all family information into 9 parsimonious indices, each of which reflect a different dimension of family background. Using this approach, I find that all family background dimensions are less strongly related to adoptees' income than to own-birth children's income. However, the attenuation is lowest for parental wealth, family structure, and occupation. This pattern provides suggestive evidence that income gaps by parental income and extended family outcomes mostly reflect differences that arise before birth, while gaps by parental wealth, occupation, and family structure reflect relatively stronger differences in the post birth environment.

Overall, the results are quite consistent with previous evidence from international adoptees.

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

[35]Some research suggests that children respond more to post birth factors when their initial endowments are stronger (Cunha and Heckman (2007), Muslimova et al. (2020)). If this is true, and if adopted children are more likely to face poor prenatal or postnatal conditions, then adopted children could be less responsive to their family background than the average non-adopted child.

[36]For example, training a new gradient-boosted decision tree results in a negative R^2. Even an OLS regression results in a model with many imprecisely estimated coefficients, making it difficult to interpret.

For example, Sacerdote (2007) and Holmlund et al. (2011) find that post-birth factors explain 20 to 30 percent of the intergenerational persistence in education in the United States and Sweden. Fagereng et al. (2021) attribute about half of the intergenerational persistence of wealth to post-birth factors in Norway. The results here extend this literature by decomposing a population level estimate which is measured with a much wider set of family background characteristics. Additionally, despite its central role in intergenerational mobility analyses, this is the first study with international adoptees that focuses on long-run income ranks.[37]

7 Conclusion

This paper provides a highly detailed analysis of income gaps across family backgrounds. I develop a multidimensional approach that enables me to link many family background characteristics to child income. I apply this approach to rich Dutch administrative data containing the most comprehensive family information studied to date. The results show that intergenerational mobility measures based on parental income only substantially understate the (dis)advantages children face, especially among the most disadvantaged children.

I then use a novel Shapley value decomposition to identify the family characteristics that explain most of the income gaps. Parental and extended family income and wealth account for the largest share of the disparities. However, accurate identification of the most disadvantaged children requires more information. Their low expected income reflects the cumulative negative contribution of many different family characteristics.

Two extensions illustrate the wider applicability of income predictions based on many family background characteristics. The first extension uses these predictions to produce new estimates of neighborhood upward mobility that adjust for differences in many observable family characteristics rather than parental income alone. The second extension applies the

[37] Sacerdote (2007) also considers international adoptees and examines child income. However, as the author acknowledges, the income measure is imperfect, complicating comparability with the broader population. See Black et al. (2020) for an analysis with domestic adoptees.

predictive model to international adoptees to quantify how much of the observed income gaps is due to differences in post-birth factors.

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 across several domains of intergenerational mobility research. As more comprehensive data becomes 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.

Black, Sandra E, Paul J Devereux, Petter Lundborg, and Kaveh Majlesi. 2020. "Poor Little Rich Kids? The Role of Nature versus Nurture in Wealth and Other Economic Outcomes and Behaviours." *The Review of Economic Studies* 87 (4): 1683–1725.

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. Forthcoming. "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility." *American Economic Review.*

Chetty, Raj, and Nathaniel Hendren. 2018. "The Impacts of Neighborhoods on Intergenerational 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.

Chyn, Eric, and Lawrence F Katz. 2021. "Neighborhoods matter: Assessing the evidence for place effects." *Journal of Economic Perspectives* 35 (4): 197–222.

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.

Cunha, Flavio, and James Heckman. 2007. "The Technology of Skill Formation." *American Economic Review* 97 (2): 31–47.

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. Forthcoming. "The Impact of the Level and Timing of Parental Resources on Child Development and Intergenerational Mobility." *Journal of Labor Economics*.

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.

Mendolia, Silvia, and Peter Siminski. 2017. "Is Education the Mechanism through Which Family Background Affects Economic Outcomes? A Generalised Approach to Mediation Analysis." *Economics of Education Review* 59 1–12.

Mogstad, Magne, and Gaute Torsvik. 2023. "Family Background, Neighborhoods, and

Intergenerational Mobility.” In *Handbook of the Economics of the Family*, edited by Lundberg, Shelly, and Alessandra Voena Volume 1. 327–387.

Muslimova, Dilnoza, Hans van Kippersluis, Cornelius A. Rietveld, Stephanie von Hinke, and Fleur Meddens. 2020. “Dynamic Complementarity in Skill Production: Evidence From Genetic Endowments and Birth Order.” Tinbergen Institute Discussion Paper 2020-082/V.

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. 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					
Year of birth	1984.6	2.9			0
Male	0.51	0.50			0
Family size	2.7	1.3			0
Household income	102156	65404			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			0.009
Primary income share	0.794	0.268			0.011
Highest education	12.937	3.637			0.358
Total wealth rank	0.50	0.29			0.008
Bank and savings balances	52,249	180,945			0.008
Bonds and shares	36,704	347,226			0.008
House value	309,747	379,964			0.008
Entrepreneurial assets	15,028	132,290			0.008
Other real estate	30,253	277,509			0.008
Substantial interest	65,601	1,235,768			0.008
Other assets	6,091	111,069			0.008
Total debt	159,239	374,080			0.007
Mortgage debt	134,709	190,726			0.008
Relationship status of household head(s) of child at age 15:					
Registered partners	0.824	0.381			0.023
Non-registered partners	0.037	0.19			0.023
Single parent	0.126	0.332			0.023
Other	0.012	0.11			0.023
Other family characteristics	**Father**		**Mother**		
Personal income	68,129	51,443	29,157	21,734	0.108
Personal earnings	83,082	61,812	33,161	26,958	0.180
Most important source of income					
Employment	0.669	0.416	0.536	0.433	0.055
Bonds or shares	0.043	0.179	0.012	0.090	0.055
Entrepreneurship	0.116	0.288	0.066	0.218	0.055
Substantial interest	0.005	0.051	0.03	0.123	0.055
Unemployment benefits	0.025	0.091	0.017	0.062	0.055

Welfare benefits	0.022	0.132	0.046	0.187	0.055
Other social security	0.004	0.049	0.007	0.062	0.055
Disability insurance transfers	0.079	0.237	0.064	0.212	0.055
Pension	0.023	0.109	0.037	0.147	0.055
Other	0.014	0.087	0.185	0.338	0.055
Type of housing					
Own house	0.745	0.409	0.7	0.428	0.066
Rental	0.053	0.19	0.104	0.259	0.066
Subsidized rental	0.2	0.356	0.195	0.338	0.066
Years of education	12.785	3.832	11.934	3.666	0.53
Average hourly wage	32.005	26.927	20.691	18.097	0.315
Most important sector of employment	In 68 categories				0.315
Suspected of any crime	0.067	0.25	0.023	0.15	0.014
Suspected of property crime	0.014	0.119	0.008	0.09	0.014
Suspected of violent crime	0.025	0.156	0.006	0.079	0.014
Suspected of other crime	0.042	0.2	0.012	0.11	0.014
Total health costs	2,700	7,153	2,626	8,212	0.014
General practitioner costs	174	143	197	155	0.063
Mental health care costs	234	3,541	321	3,948	0.063
Hospital care costs	1,830	6,723	1,692	5,013	0.063
Pharmaceutical care costs	527	2,230	542	2,084	0.063
Dental care costs	46	303	44	299	0.063
Age at first birth	29.285	5.546	26.952	4.394	0
Family size	4.14	2.365	4.044	2.299	0.218
Birth order	2.481	1.777	2.502	1.8	0.218
Father/mother not identified	0.025	0.157	0.002	0.049	0
Father/mother dead	0.008	0.086	0.004	0.065	0.019
Father/mother present in household	0.857	0.35	0.962	0.191	0.037
Migration background	In 8 categories				0.315
Migration background grandfather	In 8 categories				0.315
Migration background grandmother	In 8 categories				0.315
Extended family outcomes					
Average income rank	0.496	0.222	0.495	0.224	0.246
Average education	12.61	3.155	12.732	3.103	0.42
Average wealth rank	0.514	0.226	0.511	0.227	0.239
Average health expenditures	2717	5537	2564	5370	0.231
% of siblings suspected of any crime	0.043	0.142	0.048	0.153	0.231

Note: this Table presents descriptive statistics of the income sample. The sample comprises of all $n = 1,703,038$ children born between 1980 and 1989 with non-missing income (96.6%). 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.[38] 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.

[38] One 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

	Predicted Income Bins								
	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	18.28	21.28	25.78	31.04	50.53	65.82	69.58	73.29	77.58
	Family background characteristics								
Parental income rank	6.32	8.33	11.88	16.28	49.22	87.51	93.24	97.04	98.44
Parental wealth rank	12.11	12.55	14.16	17.32	50.86	74.13	80.33	86.65	89.65
Max. education parents	8.16	8.78	9.52	9.86	13.08	16.10	16.69	17.29	17.44
Health costs parents	5,402	5,136	4,185	3,909	2,596	1,886	1,807	1,711	1,544
Crime father	0.59	0.46	0.32	0.18	0.05	0.02	0.02	0.03	0.03
Extended family income	17.01	20.91	25.46	30.38	49.10	64.5	69.4	74.73	79.46
Extended family wealth	21.96	23.95	26.8	30.92	51.05	63.71	67.82	70.99	74.09
Father presence	0.36	0.34	0.44	0.62	0.88	0.97	0.98	0.98	0.98
Migration background	0.30	0.38	0.49	0.52	0.18	0.10	0.12	0.14	0.15
Age at first birth mother	21.80	22.64	24.11	25.31	27.05	28.39	28.64	28.90	29.00
N	1,703	1,703	13,624	17,030	272,487	17,030	13,624	1,703	1,704

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	0.025% lower bound	97.5% upper bound
(1)	(2)	(3)	(4)	(5)
0.01	3,406	0.139	0.118	0.163
0.02	6,812	0.148	0.132	0.166
0.05	17,031	0.153	0.143	0.162
0.1	34,061	0.159	0.152	0.166
0.2	68,122	0.159	0.154	0.164
0.4	136,243	0.164	0.160	0.167
0.6	204,365	0.164	0.160	0.166
0.8	272,486	0.163	0.161	0.166

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, 4, and 5 report the R^2 and 95% confidence interval lower and upper bounds, respectively. Each model is trained on a randomly selected 80% of the respective sample, and evaluated on the remaining 20%. Confidence intervals for the R^2 are bootstrapped from the test-data using 599 draws.

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

	R^2	0.025% lower bound	97.5% upper bound
Years of income	A. Varying years of income measurement		
1	0.138	0.134	0.142
2	0.145	0.142	0.150
3	0.151	0.147	0.156
4	0.153	0.149	0.157
5	0.157	0.153	0.162
6	0.158	0.154	0.162
7	0.162	0.158	0.166
8	0.161	0.157	0.166
9	0.165	0.161	0.170
All	0.170	0.166	0.174
All > age 32	0.166	0.162	0.170
Age child	B. Varying ages of income measurement		
30-33	0.129	0.125	0.133
34-37	0.154	0.150	0.159
38-41	0.153	0.149	0.158

Notes: each row presents the R^2 and corresponding 95% lower and upper bound for 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. Confidence intervals for the R^2 are bootstrapped from the test-data using 599 draws.

Table A5: Within Neighborhood Estimates

	Household income rank	
	(1)	(2)
Predicted income rank	1.001	0.964
	(0.002)	(0.003)
Neighborhood Fixed Effects		x
N	1,655,052	1,655,052
R^2	16.5%	17.0%

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)	0.045
Signal SD(MAUM)	0.020
	B. Estimation with a second order polynomial
Signal SD(AUM)	0.048
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)	0.051
Signal SD(MAUM)	0.024
	E. Estimation at the 75th percentile of the parental/predicted income distribution
Signal SD(AUM)	0.033
Signal SD(MAUM)	0.019
	F. Relative intergenerational mobility
Signal SD($\hat{\beta}_n$)	0.150
Signal SD($\hat{\gamma}_n$)	0.072
	G. (Multidimensional) Absolute Upward Mobility estimates at the municipality level
Signal SD(AUM)	0.030
Signal SD(MAUM)	0.019

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	0.677
	(0.004)
N	165,256

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: Descriptive Statistics for International Adoptees and their Parents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income rank (y)	35.58	37.45	39.16	38.61	42.09	43.24	41.17	41.36	43.44	42.97
Predicted income rank ($\hat{y}$)	38.16	45.63	49.13	51.84	54.27	56.56	58.69	60.95	63.89	69.38
	Characteristics Adoptive Parents									
Parental income rank	19.24	29.74	37.08	43.22	52.84	59.83	67.7	76.23	84.71	93.44
Parental wealth rank	32.33	46.47	54.16	58.51	62.86	66.55	69.12	70.29	72.33	81.53
Highest education parents	11.18	11.99	12.89	13.27	14.42	14.6	14.74	15.42	15.87	16.42
Father suspected of crime	0.1	0.04	0.04	0.03	0.04	0.02	0.03	0.02	0.02	0.02
Health expenditures parents	4,386	3,732	3,444	2,829	3,289	2,594	2,696	2,638	2,707	2,228
Extended family income rank	37.97	44.35	47.34	50.85	51.61	55.93	58.5	60.6	62.12	70.99
N	504	504	505	504	505	504	504	505	504	505

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 A9: The Effect of Family Background on Income: Regression Results with Adoptees

	Household income rank			
	(1)	(2)	(3)	(4)
Predicted income rank	0.279	0.283	0.284	0.273
	(0.031)	(0.031)	(0.032)	(0.028)
Age at migration control		x	x	x
Gender FE		x	x	x
× Country of Origin FE			x	x
× Year of Adoption FE				x
N	5,044	5,044	5,044	5,044

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 A10: Regression Results with Adoptees: Heterogeneity by Origin Country

	Household income rank					
	(1)	(2)	(3)	(4)	(5)	(6)
Predicted income rank	0.223***	0.348***	0.346***	0.237	0.393**	0.362*
	(0.072)	(0.100)	(0.127)	(0.147)	(0.196)	(0.197)
Origin country	Sri Lanka	Indonesia	South Korea	Colombia	India	Brazil
N	2,029	957	624	529	292	250

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. ($^{***}: p < 0.01, ^{**}: p < 0.05, ^{*}: p < 0.1$)

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

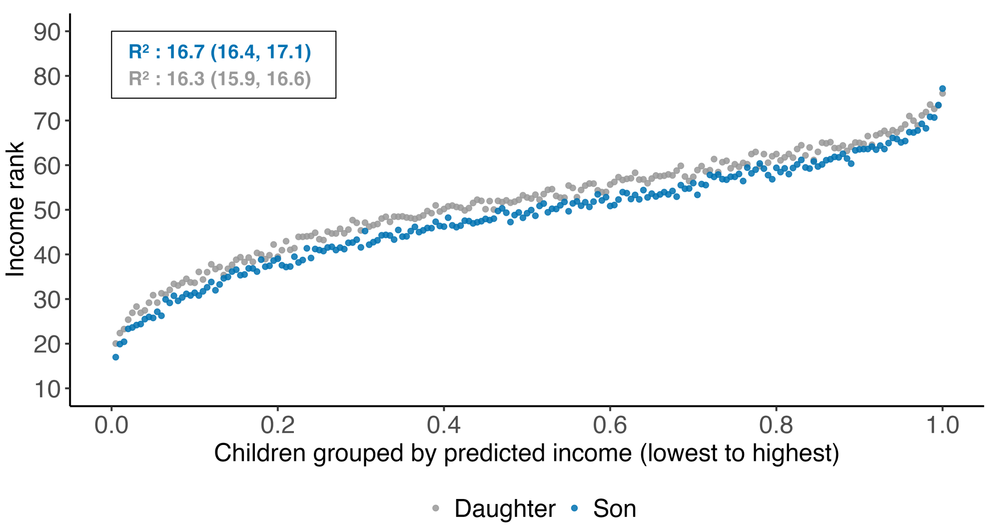

Notes: this Figure presents binscatter plots of sons' and daughters' household income ranks for 173,652 sons and 166,957 daughters in the test data, who are sorted into bins based on their predicted income rank. Predictions are generated using the same predictive model and explanatory variables as in 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

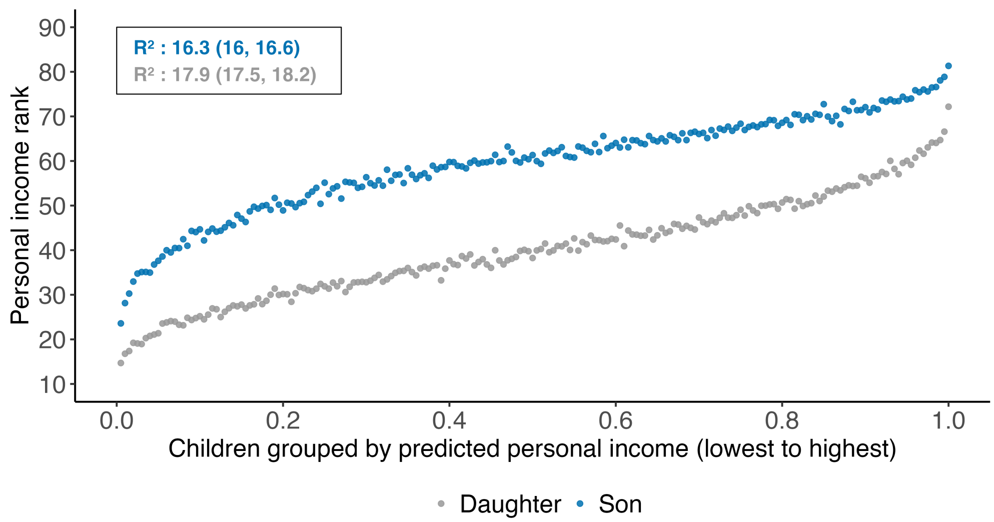

Notes: this Figure presents binscatter plots of sons' and daughters' personal income ranks for 172,976 sons and 164,990 daughters in the test data, who are sorted into bins based on their predicted income rank. The graphs are constructed using the same steps as in Figure 2, applied to children's personal 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 (Disposable) Household Income Levels

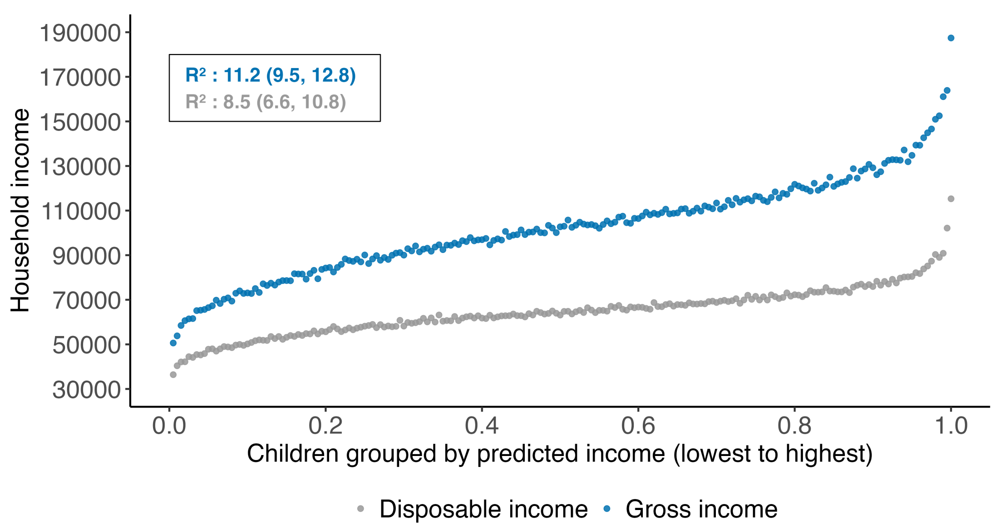

Notes: this Figure presents binscatter plots of children's gross household income and disposable household income for 340,608 children in the test data, who are sorted into bins based on their predicted income rank. The graphs are constructed using the same steps and sample as in Figure 2, applied to children's gross household income and disposable household income levels instead of ranks. Confidence intervals for the R^2 are bootstrapped from the test data using 599 samples and are reported in brackets

Figure A4: Predicting Children's Education and Crime

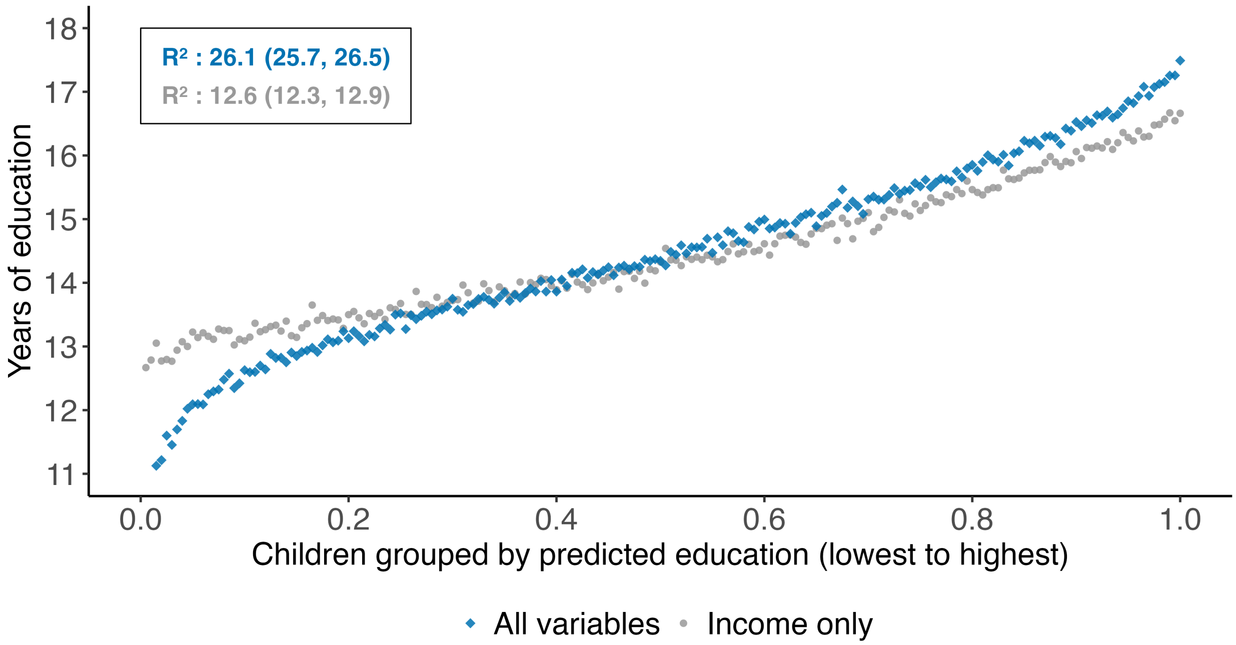

(a) Education

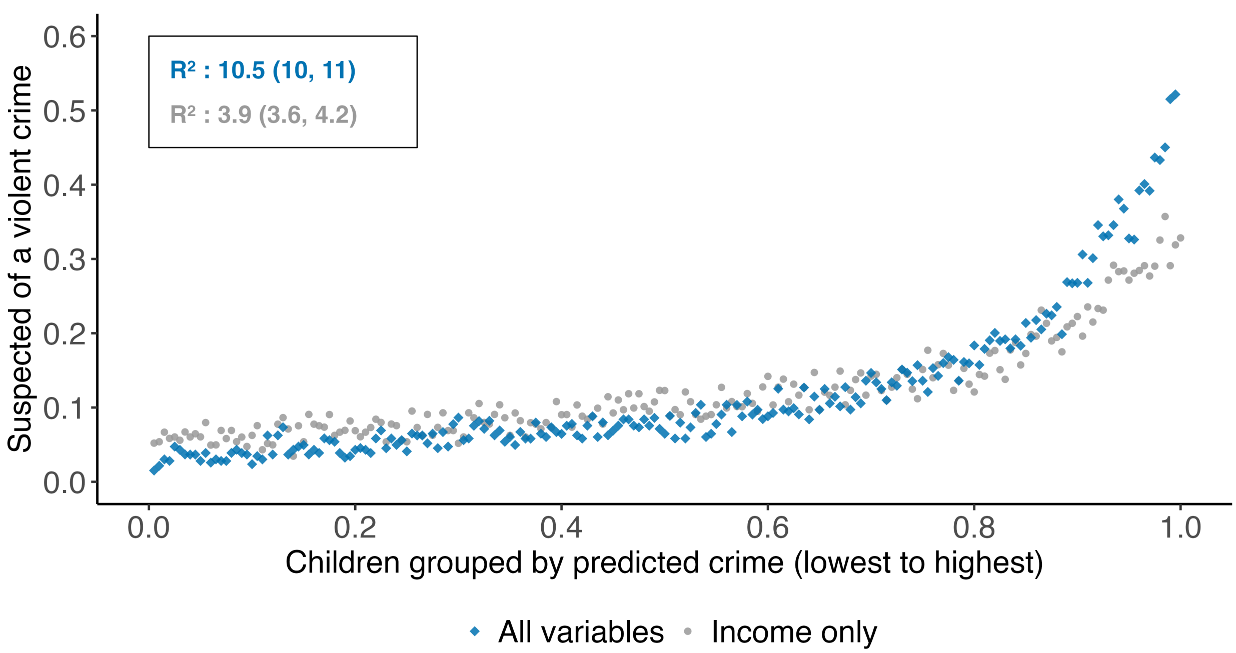

(b) Crime (sons only)

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 180,829 children from the test sample. Panel B reports the results for 92,725 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 A5: Comparing Neighborhood Upward Mobility Rankings

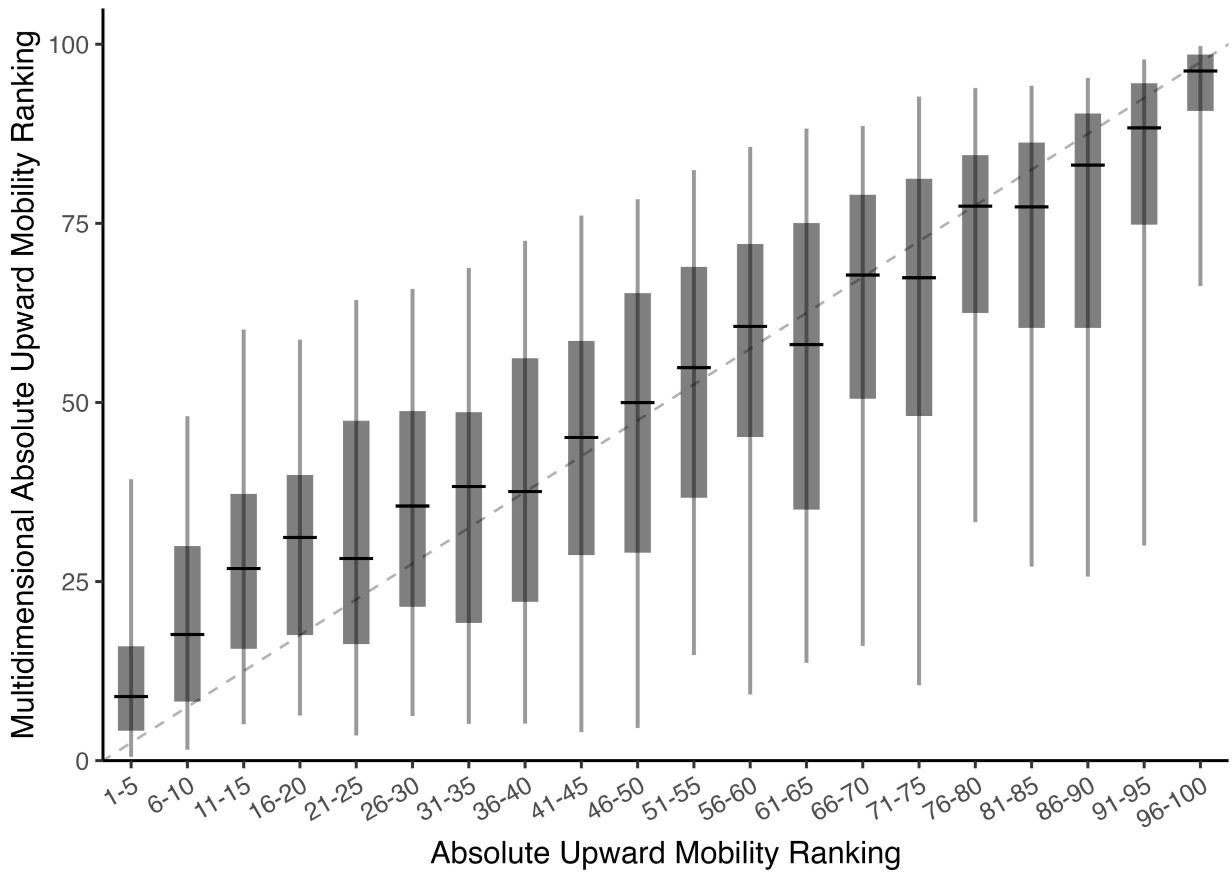

Notes: this Figure compares neighborhood rankings based on Absolute Upward Mobility with rankings based on Multidimensional Absolute Upward Mobility. All 2,828 neighborhoods are ranked from 0 to 100 on each measure based on the estimates in Figure 5. Neighborhoods are grouped into twenty equally sized 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 that 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, coincidentally, equal up to the second digit. 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 31% of all variation in income ranks is driven by factors shared between siblings.

Figure B1 reports a 5 $\times$ 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	0.323	0.324	0.288	0.290	-
	(0.001)	(0.001)	(0.001)	(0.001)	
N	1,702,355	1,702,355	864,064	825,170	1,702,355
Adjusted R^2	0.105	0.091	0.093	0.095	0.308

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

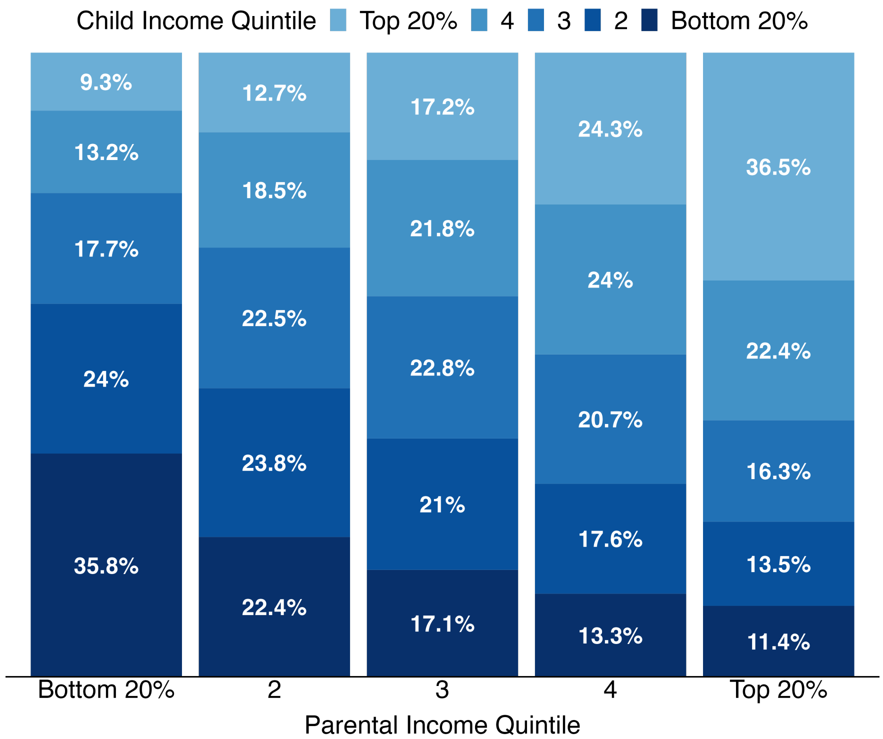

Notes: this Figure presents the transition matrix of child income conditional on parental income quintile. Each bar represents the distribution of child income quintiles for children whose parents fall in the corresponding parental income quintile on the x-axis. The segments within each bar show the 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 0.32 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 2003 and 2013. 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	0.297	0.304	0.311	0.316	0.320	0.323	0.325	0.327	0.329
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
N	1,098,025	1,098,025	1,098,025	1,098,025	1,098,025	1,098,025	1,098,025	1,098,025	1,098,025
R^2	0.082	0.086	0.090	0.092	0.095	0.096	0.098	0.099	0.100

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	0.290	0.294	0.292
	(0.001)	(0.001)	(0.001)
Years of income measurement parents	2003-2007	2006-2010	2009-2013
N	1,267,606	1,267,606	1,267,606

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 2003 and 2013. Standard errors are reported in parentheses.

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

	(1)	(2)	(3)
Coefficient	0.274	0.304	0.308
	(0.002)	(0.002)	(0.002)
Age child	30-33	34-37	38-41
N	326,388	326,388	326,388

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 3.4% 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 0.023 (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 2023. Implementing their age at measurement further reduces the rank-rank correlation by 0.021 (column 4).

Third, Van Elk et al. measure parental income from 2003 to 2005. I measure parental income over all available observations from 2003 and up to age 60. On average, that corresponds to 12 observations for fathers and 14 observations for mothers. Implementing their parental age at measurement further reduces the estimate by 0.029 (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 0.055, 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.[39]

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

	(1)	(2)	(3)	(4)	(5)	6)
Coefficient	0.329	0.313	0.306	0.285	0.256	0.201
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
N	1,016,358	883,471	779,159	779,159	778,998	778,998
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 2003 (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 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 2003 and up to age 60.

[39]Using 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.

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 2003 to 2009. While I cannot observe incomes prior to 2003, this at least aligns the number of years over which parental income is measured.[40] This reduces the rank–rank correlation to 0.29. 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 with Manduca et al. is that Boustan 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 0.22.

Since I do not observe parents' incomes before 2003, 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.

[40] Estimates are stable across different years of income measurement between 2003 and 2013 (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	0.327	0.292	0.256	0.229
	(0.001)	(0.001)	(0.001)	(0.001)
N	986,125	986,125	986,125	986,125
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 2003) 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

This section studies which adoptive family characteristics are most strongly associated with adoptees' income. This is challenging because the ratio of explanatory variables to the number of observations is much larger than in the main results section. In such settings, nonparametric predictors like decision trees can perform very poorly compared to simpler models that impose stricter functional form assumptions. In fact, estimating a new gradient-boosted decision tree using the sample of international adoptees results in a negative R^2. While a simple OLS regression where all variables enter linearly results in slightly higher adjusted R^2 of 1.4 percent, it has many imprecisely estimated coefficients, rendering it difficult to interpret.

To improve interpretability, I summarize all family information into a few highly parsimonious indices. I do this using the Shapley values of the income predictions from the comprehensive model in Section 4. For each child, I first compute the Shapley values of all the family background characteristics. I then sum the Shapley values of all family background characteristics that belong to the same category in Table 1. This produces nine indices that each measure the total contribution of one family background dimension to a child's predicted income. For example, the wealth index equals the summed contribution of all wealth related predictors to a given income prediction. By construction, the sum of these indices equals the total prediction for each child (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 these indices for both adoptees and own-birth children from the test sample. Table C1 column 1 shows the results for the own-birth children. To guide interpretation, consider the coefficient of 0.059 for the parental income index. It implies that children from families who are one standard deviation more advantaged in terms of the parental income index, but are similarly advantaged in all other dimensions, have income that is 5.9 ranks higher on average. Each index is constructed so that higher values correspond to higher expected child income. For instance, a higher parental crime index does not indicate that parents were more likely suspected of crime. It indicates a lower chance of parental crime and therefore a higher expected income for their children. This design ensures that all coefficients have the same sign and are easy to compare.

Consistent with the main results, the joint explanatory power of the indices is 16.6 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.

The explanatory power of the same regression for adoptees is 1.2 percent. This is close to the adjusted R^2 of an OLS model that includes all variables separately, which is 1.4 percent. This shows that the parsimonious model, despite its limited degrees of freedom, captures almost the same share of income variation as the full set of family background variables. The main advantage is that the coefficients in the parsimonious model are estimated with much greater precision, which makes them easier to interpret.

The associations between the indices and adoptees' income are much weaker than those for own birth children. The coefficient for the parental income index, for example, is about one-sixth of that for own-birth children. It implies that being raised in an adoptive family which is one standard deviation ‘more advantaged’ in terms op parental income, but similarly advantaged in all the other dimensions, raises adoptees’ income by 1 rank. The only index that does not show attenuation is the one for parental crime, but its estimate is imprecise because crime among adoptive parents is relatively rare. The 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 only marginally significant, so this result should be interpreted with caution.

Interestingly, the attenuation is smaller for parental wealth, occupation, and family structure. A one standard deviation increase in the parental wealth or family structure index has even higher effects than a one standard deviation increase in the parental income index. This stands in contrast to the results for own birth children. A possible explanation is that parental income, education, and extended family outcomes correlate more strongly with genetic endowments that parents pass to their biological children. These channels are absent for adoptees, which reduces the predictive power of these indices. If instead parental wealth, occupation, and family structure are better proxies of the environment in which adopted children grow up, then their predictive power falls less when the biological link is removed.

To sum it up, linking adoptees’ income to many family background characteristics without strong functional form assumptions remains difficult. Combining all family information into indices improves interpretability with little loss in explanatory power. Parental wealth, family structure, and occupation are the strongest predictors of adoptees’ income.

Table C1: The Relationship Between all Family Characteristics and Child Income

	Household income rank	
Parental income	0.059***	0.010*
	(0.001)	(0.005)
Parental wealth	0.032***	0.014***
	(0.001)	(0.005)
Extended family outcomes	0.026***	0.006
	(0.001)	(0.005)
Family structure	0.021***	0.014***
	(0.000)	(0.005)
Parental occupation	0.015***	0.010**
	(0.001)	(0.005)
Migration background	0.009***	-0.014*
	(0.000)	(0.007)
Education	0.008***	0.001
	(0.001)	(0.005)
Parental health expenditures	0.007***	0.002
	(0.000)	(0.004)
Parental crime	0.005***	0.009
	(0.000)	(0.006)
N	340,608	5,044
R^2	16.6%	1.2%
Sample	Own-birth children	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 column 2 corresponds to the adoption sample, as in Figure 6. The construction of the indices is explained in the text above. Standard errors are reported in parentheses. ($^{***}: p < 0.01$, $^{**}: p < 0.05$, $^{*}: p < 0.1$)