TI 2024-063/V
Tinbergen Institute Discussion Paper

The causal component in the intergenerational transmission of income

Monique de Haan[1]
Magnus Stubhaug[2]

1 Amsterdam School of Economics, University of Amsterdam
2 NIFU, Oslo.

The causal component in the intergenerational transmission of income

Monique De Haan* Magnus Stubhaug†

October 2024

Abstract

We apply a partial identification analysis using comprehensive Norwegian register data to investigate the causal effect of father's income on child income. We find a strong association between the incomes of fathers and children. The causal effect, however, equals at least 1% and at most 51% of this observed association. Additionally, we find substantial differences in the intergenerational association and bounds around the causal effect between sons and daughters when considering individual incomes. When examining joint income with their partners, the results are more aligned, indicating that assortative mating plays a key role in intergenerational income transmission, particularly for daughters.

Keywords: Intergenerational income transmission, partial identification, intergenerational mobility
JEL Classification: J62, C14, D31

*Amsterdam School of Economics, University of Amsterdam. Also affiliated with the Tinbergen Institute, CESifo, Norwegian Fiscal Studies (NFS). Corresponding author: Email: monique.dehaan@uva.nl

†NIFU, Oslo.

1 Introduction

A sizable and expanding literature documents a strong correlation between parental incomes and those of their offspring.[1] This literature primarily centers on measuring the intergenerational persistence in incomes and illustrates how this persistence varies between countries and over time. Our understanding of the determinants of intergenerational income transmission remains however incomplete. As noted by Mogstad and Torsvik (2023), while it is crucial to measure intergenerational persistence, our ultimate objective is to understand why family background exerts such a significant influence on income and other life outcomes.Three primary reasons have been proposed in the literature to explain the observed positive association between the incomes of parents and children: (1) Innate abilities leading to high incomes are genetically transmitted from parents to children; (2) High-income parents possess different child-rearing skills compared to low-income parents, positively influencing their children's incomes; (3) Parental income causally affects child income, potentially through monetary investments in their children's human capital. This paper focuses on the third reason and estimates the portion of the intergenerational income association reflecting a causal effect.

Knowledge about the magnitude of the causal effect of parental income on child income is important for, among others, policymakers. A substantial causal effect implies that redistributive policies in the current generation could alter income inequality in the next generation. Moreover, policy interventions that unintentionally redistribute incomes in the current generation may have lasting effects, impacting the incomes of future generations (e.g., policies enacted in response to the Covid-19 pandemic). To assess the extent to which the intergenerational income association is attributable to a causal effect, we combine a partial identification approach with comprehensive Norwegian register data, including 50 years of income data from the tax administration, and unique identifiers linking parents to their children.

We contribute to the existing literature in various ways. First, we contribute to the literature that measures the intergenerational association in incomes by disentangling the causal component. Secondly, we contribute to a smaller body of literature that breaks down the intergenerational income association into a nature and nurture component. Here, "nature" encapsulates genetic inheritance (reason 1), while "nurture" encompasses both the impact of parents' child-rearing skills and the potential causal effect of parental income on child income (reasons 2 and 3). Some studies in this second strand of the literature exploit the absence of a genetic link between adoptees and their adoptive parents (Sacerdote (2007); Björklund et al. (2006)). Conversely, Amin et al. (2011) adopt a different strategy, exploiting that (monozygotic) twins have the same genetic makeup, and control for genetic inheritance by including twin-parent fixed effects in the regression of son's income on father's income. We contribute to these studies by decomposing the causal component rather than the nurture component, which also includes the effect of parenting styles and other child-rearing skills. Additionally, we concentrate on the entire population rather

[1] see Nybom (2024) and Mogstad and Torsvik (2023) for recent overviews of this literature

than specific subpopulations of twins or adoptees.

Another study that decomposes the intergenerational income elasticity is Lefgren et al. (2012). They use Swedish data and combine linearity and homogeneous treatment effect assumptions with assumptions about the direct effect of different instrumental variables on potential incomes of children to decompose the intergenerational elasticity. Lefgren et al. (2012) attribute the majority of the intergenerational income elasticity to the intergenerational transmission of endowments as opposed to a causal effect of financial resources. In contrast to Lefgren et al. (2012), our empirical approach is fully nonparametric, thereby allowing both for non-linear as well as heterogeneous effects of parental income on child income.

A third strand in the literature exploits quasi-experimental variation in parental income to estimate the causal effect on child outcomes. As discussed in the overview of this literature by Cooper and Stewart (2021), most of the studies exploit variation at the lower end of the parental income distribution, such as cash transfers or changes in the EICT in the U.S.. Since there are hardly any quasi-experimental studies that use children's income as outcome variable, Cooper and Stewart (2021) focus in their review on the effect of parental income on short term and intermediate term health and schooling outcomes of children and show that most studies find positive effects on these outcomes.

A recent quasi-experimental study that uses child income as outcome variable is Aizer et al. (2016). They estimate the effect of receiving transfers under the Mothers' Pension program in the U.S and find a positive effect on child income in young adulthood. In our analysis we do not exploit quasi-experimental variation in parental income, but instead use nonparametric assumptions to bound the causal effect of parental income on child income. This approach allows us to estimate the average causal effect of income changes for the full population, along the entire distribution of father's income on long-term child income.

We start our analysis by estimating the intergenerational association in income and find that also in Norway there is a strong relation between the income of fathers and their offspring. Increasing a father's income from the bottom five percent to the top five percent of the income distribution—a rise of approximately 780 percent—corresponds to an average increase of 419,931 NOK ($52,239) in the child's income, which represents a 140 percent increase on average.[2] This implies an observed intergenerational income elasticity of 0.18.

Next, we introduce assumptions to bound children's mean potential incomes as well as the causal effect of an increase in father's income. We build on the observation that children with high income fathers tend to differ in terms of education, ability and other types of skills compared to children from low income fathers. Based on this, we assume that the potential incomes of children with high-income fathers are not, on average, lower than those of children with low-income fathers (monotone treatment selection). In addition, we assume that increasing father's income will, on average, not reduce the adult incomes of his offspring (monotone treatment

[2]When we covert amounts to dollars we use the average Norwegian Krone to US Dollar exchange rate in 2015 (1 Norwegian Krone = 0.1244 US Dollar.) Source: https://www.exchange-rates.org/exchange-rate-history/nok-usd-2015 (last accessed 17/09/2024)

response). We continue by exploiting variation in childhood neighborhood income and parental education to tighten the bounds. We first use these background characteristics as conditioning variables and assume that the monotone treatment selection and monotone treatment response assumptions hold conditional on neighborhood income and parental education. Next we use both variables as monotone instruments (Manski and Pepper (2000)) and assume that children's mean potential incomes are non-decreasing in the combination of neighborhood income and parent's level of education.

The tightest bounds show that the causal effect of father's income is substantially smaller than the intergenerational association. Increasing father's income from the bottom five percent to the top five percent of the income distribution increases child income on average by at least 3,396 NOK ($422) and at most 212,464 NOK ($26,431), which means that the causal component is at least 1 percent and at most 51 percent of the corresponding intergenerational association. These results imply a causal intergenerational income elasticity of at most 0.08.

We further find that the income association between sons and their fathers is much stronger than it is for daughters. The estimated bounds on the mean potential incomes further show that, at most, there is a very small causal effect of father's income on daughter's income, while the upper bound on the causal effect for sons is three times larger. When we consider the combined income of both the father and mother instead of solely the father's income, we observe that this does not alter the observed differences between sons and daughters. However, when we change the outcome to include the joint income for children and their partners, we find that both the intergenerational associations and the bounds around the causal effects become more comparable between sons and daughters. This suggests that assortative mating plays an important role in the intergenerational transmission of income, in particular for daughters.

The remainder of the paper is organized as follows. In the next section we present the data and the construction of the income variables. We present our identification strategy in Section 3. In Section 4 we show our results and in Section 5 we summarize and conclude.

2 Data

In our analysis, we utilize income data extracted from the Norwegian tax registers spanning the years 1967 to 2017. We combine this with detailed individual-level information from the Norwegian population registers. As indicated by a number of studies, estimates of intergenerational income persistence can suffer from life-cycle bias.[3] This bias tends to be smallest when the incomes of both parents and children are measured in midlife (Nybom and Stuhler (2016)), and it is further important to use as many years as possible to construct the income variables. Our analysis is therefore centered on individuals born in Norway between 1960 and 1967, because for these cohorts we are able to establish a midlife income measure for two generations. For the child generation we do this by averaging income data over 11 years, when the children are

[3] See among others Haider and Solon (2006), Böhlmark and Lindquist (2006), and Nybom and Stuhler (2016).

aged between 40 and 50. This income data comprises yearly individual-level taxable income, encompassing all labor and capital income after deductions.[4] We next successfully match 98.9 percent of the children to their fathers using unique personal identifiers. For the fathers, we construct an income measure by averaging yearly taxable incomes when they were between 40 and 60 years old. Given that tax registers in Norway commence in 1967, and some fathers are already over 40 years old by then, we include income years up to the age of 60. Additionally, we require that we observe the father's yearly income for at least 10 out of the 21 years, leading to a 6 percent reduction in the sample size. The final sample consists of 427,512 children and their fathers. We adjust the incomes for both fathers and children to 2015 NOK to account for inflation.

As we explain in Section 3, we use a nonparametric method to estimate bounds on the causal effect of an increase in father's income on the income of his offspring. Our methodology involves categorizing the father's income into 12 distinct categories. First, we create deciles based on the distribution of father income. Subsequently, we split the top and bottom deciles in half to establish categories for the top and bottom five percent, respectively. With these categories in place, we then estimate the bounds on the causal effect of increasing father's income from one category to the next. This approach allows us to examine how both the intergenerational association in income and the causal effect vary along the distribution of father's income, while ensuring an adequate number of observations in each category.

The lower panel in Figure 1 illustrates the density of father's income, with the 12 categories marked by vertical lines. Meanwhile, the upper panel in Figure 1 shows the mean of children's observed income corresponding to each of these 12 categories of father's income. There is a strong positive relation between father's and children's incomes along the entire range of father's income. Notably, increasing father's income from the bottom five percent to the top five percent is associated with an average increase in child income of 419,931 NOK ($52,239), equivalent to an increase by on average 140 percent. If this association entirely reflects a causal effect, then redistributing income within the father's generation could have substantial consequences for the child generation. However, given that the observed association might also stem from genetic inheritance or differences in parental child-rearing abilities, the subsequent section introduces the empirical methodology we employ to disentangle the causal component.

[4]In certain years covered by the tax register, income below the tax threshold is not recorded. Although this threshold varies slightly over the years, it generally remains very low. For these instances, we assign zero income to individuals that are registered as alive and living in Norway in the population register, but who are absent from the tax register.

Figure 1. The association between father's and children's incomes

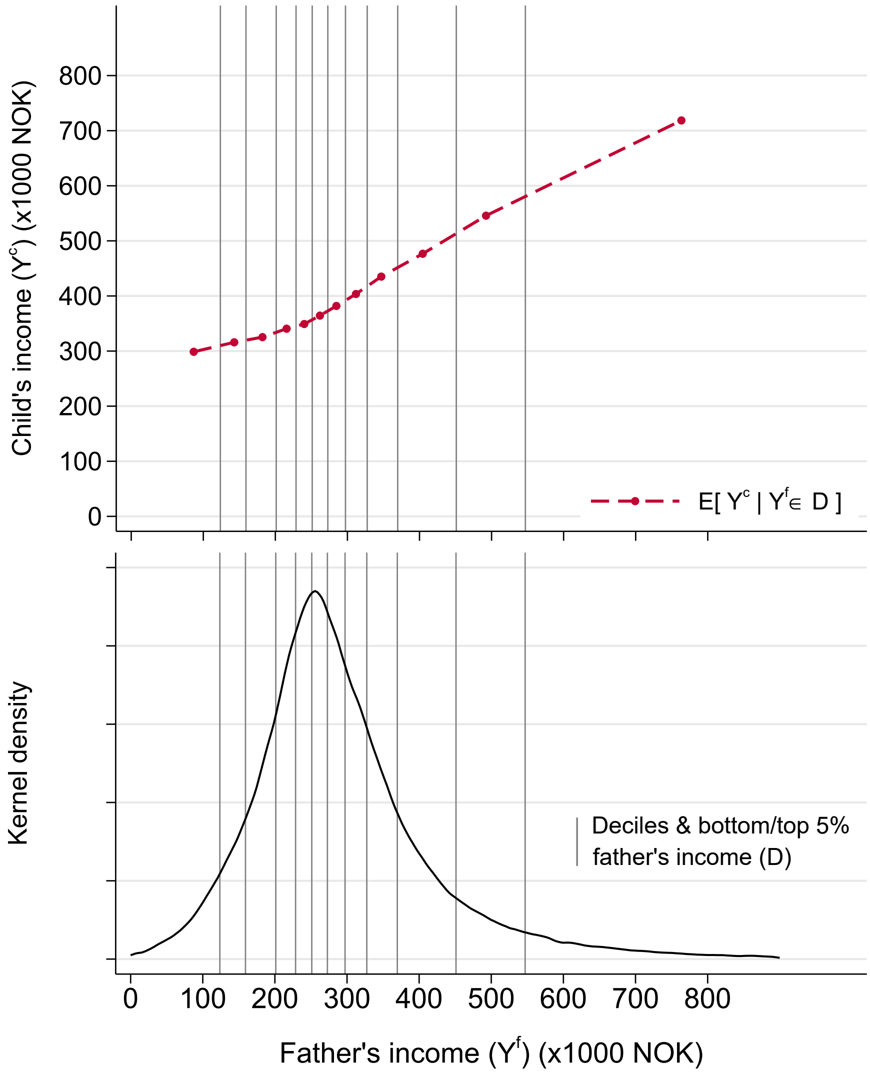

Note: The bottom panel shows the distribution of father's income (in 1000 NOK) . The vertical lines show how our treatment levels ($D \in [1,12]$) are defined. The top panel plots the observed mean income (in 1000 NOK) among children whose father has an income in the particular category. Incomes for both fathers and children are adjusted to 2015 NOK to account for inflation. Number of observations equals 427,512.

3 Empirical approach

We are interested in the average causal effect of an increase in father's income on children's income

$$E[Y^c(D^f = d^*) - Y^c(D^f = d)] = E\left[Y^c\left(d^*\right)\right] - E\left[Y^c\left(d\right)\right], \tag{1}$$

where Y^c is child's income and D^f is father's income category as defined in Section 2. $Y^c(d)$ is the child's potential income in case his or her father's income is within category d. While this potential outcome is observed for children whose father has an income in category d it is unobserved for all children for whom their father's income is in a category below or above d. This implies that the mean of the potential outcome can be written as a weighted average of observable and unobservable components

$$
\begin{aligned}
E[Y^c(d)] = \quad & E[Y^c(d)|D^f < d] \cdot Pr(D^f < d) \\
& +E[Y^c|D^f = d] \cdot Pr(D^f = d) \\
& +E[Y^c(d)|D^f > d] \cdot Pr(D^f > d)
\end{aligned} \tag{2}
$$

In the coming subsections we will show how we can use mean-monotonicity assumptions to obtain upper and lower bounds on the unobserved counterfactual mean potential outcomes $E[Y^c(d)|D^f < d]$ and $E[Y^c(d)|D^f > d]$ and on the mean potential outcome in the full population $E[Y^c(d)]$.

3.1 Monotone treatment selection

We start by assuming monotone treatment selection (MTS) (Manski and Pepper (2000)) which states that

$$E[Y^c(d)|D^f = d_2] \geq E[Y^c(d)|D^f = d_1] \quad \forall \quad d, \; d_2 > d_1. \tag{3}$$

If we compare children from high income fathers ($D^f = d_2$) to children from low income fathers ($D^f = d_1$) and consider the hypothetical situation in which all fathers have an income within category d, the MTS assumption states that the mean income we would observe for the children of initially low income fathers would not be higher than the mean income of children with initially high income fathers. This assumption is motivated by the observation that high income fathers tend to differ in characteristics from low income fathers; they are on average higher educated, but likely also differ in unobserved characteristics such as ability, motivation and preferences. Since these characteristics can be transmitted from fathers to children via genetic inheritance or via the way fathers raise their children, children of high income fathers will on average likely have a higher income themselves compared to children from low income fathers, regardless of the actual income of their father.

The MTS assumption could be violated if the traits that make fathers have a high income have the opposite effect on their children's income. For example, fathers with high incomes might earn more because they prefer to work long hours, leaving less time to spend with their children. If this negatively impacts the future earnings of their children, it could potentially lead to a violation of the MTS assumption. However, this would only pose an issue if it holds true on average, not just for some fathers. The literature does not support this concern; instead, higher earnings (potential) is linked to more time spent with children, not less (Guryan et al., 2008). Additionally, if the preference for working long hours is genetically transmitted, and the children of high-income fathers also earn high incomes due to their work ethic, this would be consistent with the MTS assumption.

Under the MTS assumption we can use the observed mean income of children whose father has an income in category d, $E\left[Y^c|D^f = d\right]$, as an upper bound on the unobserved mean potential outcome for the children whose father has an income in a lower category $E[Y^c(d)|D^f < d]$. In a similar way we can use $E\left[Y^c|D^f=d\right]$ as a lower bound on the unobserved mean potential outcome for the children whose father has an income in a higher category $E[Y^c(d)|D^f>d]$.

3.2 *Monotone treatment response*

The second assumption that we use to tighten the bounds around the mean potential outcomes is the so-called monotone treatment response (MTR) assumption (Manski (1997)). The MTR assumption states the following

$$E[Y^c(d_2)|D^f=d]\geq E[Y^c(d_1)|D^f=d] \quad \forall \quad d,\ d_2>d_1$$

If we take the sub-population of children whose father has an income in category d_1 and we consider the hypothetical situation in which their fathers' income is increased to category d_2, the MTR assumption states that the incomes of these children will not be reduced on average due to the increase in their fathers' income. It is important to note that the MTR assumption does not impose a positive causal effect but it assumes the effect to be non-negative, on average. This assumption is motivated by the theoretical literature that predicts that an increase in parental income leads to higher investments in the human capital of children (Becker and Tomes (1986)) which is in accordance with a number of empirical studies that find a positive effect of parental income on a range of child outcomes, such as health and education (Cooper and Stewart (2021)).

The MTR assumption could be violated if an increase in a father's income decreases the incentives of children to earn a high income themselves. While this might be true for some individuals, it would only violate the MTR assumption if it holds true on average across the population, which we consider unlikely. Another reason we believe this scenario is unlikely to violate the MTR assumption is that our income measure includes not only earnings but all types of income, including capital income.

Under the MTR assumption we can use the observed mean income of children whose father has an income in a category higher than d, $E\left[Y^c|D^f>d\right]$, as an upper bound on the unobserved mean potential outcome for these children in case their father would have an income in category d, $E[Y^c(d)|D^f>d]$. In a similar way we can use the observed mean income of children whose father has an income in a category lower than d, $E\left[Y^c|D^f<d\right]$, as a lower bound on the unobserved mean potential outcome for these children in case their father would have an income in category d, $E[Y^c(d)|D^f<d]$.

3.3 *Neighborhood income as monotone instrument*

Many studies show that children's outcomes are strongly related to characteristics of the neighborhood in which they grew up (Durlauf (2004)). These correlations might reflect a causal neighborhood effect, or that individuals living in different neighborhoods differ in characteristics that affect their outcomes. Unlike a recent and growing literature[5] we do not aim to identify a

[5]See for example: Kling et al. (2007); Ludwig et al. (2013); Chetty et al. (2016); Chetty and Hendren (2018).

causal neighborhood effect, but instead exploit variation across neighborhoods to obtain tighter bounds around the causal effect of father's income on his offspring's income. In particular we use average neighborhood income during childhood as a monotone instrumental variable (MIV) and assume that children's mean potential incomes are nondecreasing in neighborhood income (Manski and Pepper (2000))

$$E[Y^c(d)|N = n_2] \geq E[Y^c(d)|N = n_1] \quad \forall \quad d, \; n_2 > n_1 \tag{4}$$

If we would compare children that grew up in a high income neighborhood ($N = n_2$) to children that grew up in a low income neighborhood ($N = n_1$) and consider the hypothetical situation in which all fathers have an income within category d, the MIV assumption states that the mean income we would observe for the children from low income neighborhoods would not be higher than the mean income of children from high income neighborhoods. The MIV assumption does not impose that there is an effect of neighborhood income on children's potential incomes, but if there is an effect this is assumed to be weakly positive.

The measure of neighborhood income that we use to construct the monotone instrument is the mean taxable income among prime-aged individuals in the municipality where the child lived at age 16. We take a 3-year average to get a more precise measure and next divide the neighborhood incomes into 10 categories; 8 intervals of 10,000 NOK, and one bottom and one top category. Figure 2 shows a histogram of average neighborhood income as well as the cut-offs of the 10 categories of the monotone instrument displayed by the vertical bars.[6]

Figure 2. Histogram of neighborhood income

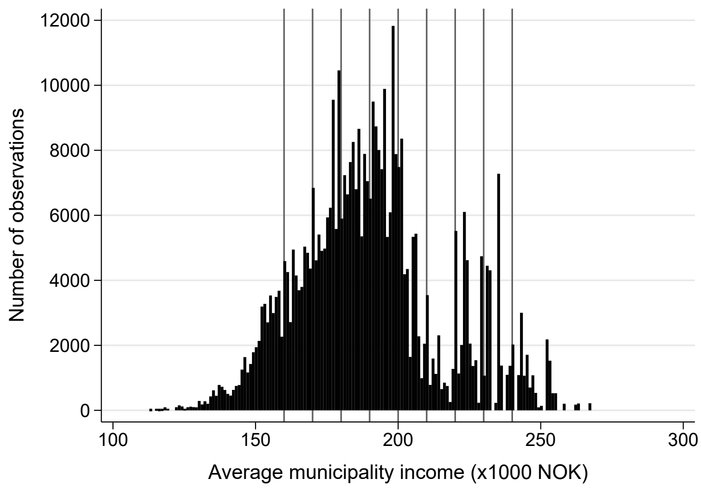

Note: The black bars show the histogram of mean taxable income among prime-aged individuals in the municipality where the child lived at age 16, with the number of observations in each 1000 NOK bin. We take a 3-year average to get a more precise measure and incomes are adjusted to 2015 NOK to account for inflation. The gray vertical lines show how we have defined the neighborhood income MIV. Total number of observations equals 427,512.

[6] In Section 4.2 we show the main results for different numbers of categories of the neighborhood income MIV.

3.4 Parent's education as monotone instrument

In addition to the variation across neighborhoods, we exploit variation across parental education groups, by using the maximum of mother's and father's levels of education as monotone instrument. Under the MIV assumption, shown in equation 5, children's mean potential income is nondecreasing in their parent's education.

$$E[Y^c(d)|E = e_2] \geq E[Y^c(d)|E = e_1] \quad \forall \quad d,\ e_2 > e_1 \tag{5}$$

If we would compare children with high educated parents ($E = e_2$) to children with low educated parents ($E = e_1$) and consider the hypothetical situation in which all fathers have an income within category d, the MIV assumption states that the mean income we would observe for the children with low educated parents would not be higher than the mean income of children with high educated parents. This assumption is motivated both by the theoretical literature on the intergenerational transmission of human capital (see e.g. Becker and Tomes (1986)) as well as the empirical literature documenting a strong relation between parental education and child outcomes (Holmlund et al. (2011)). Figure 3 shows the levels of the parental education MIV along with the number of observations in each category.

Figure 3. Histogram of parent's education

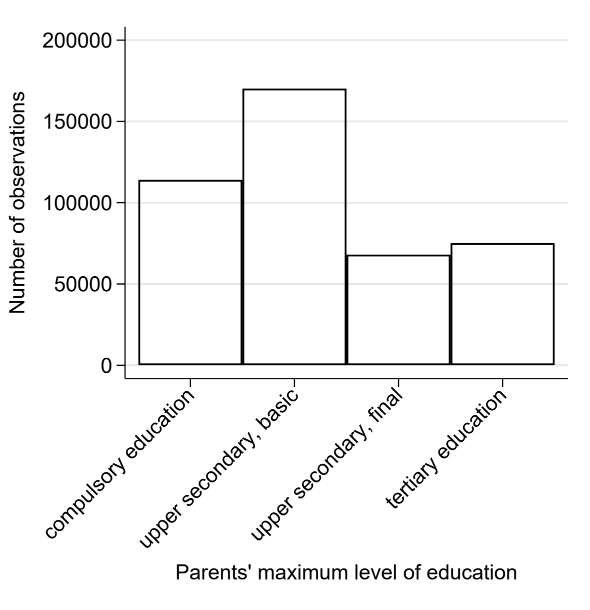

Note: The bars show the number of observations in each category of the parental education MIV. Total number of observations equals 427,512.

3.5 Combining assumptions

In Section 4 we present bounds around the mean potential incomes and the causal effect of father's income on child income by combining the above discussed assumptions. We start by combining the MTR and MTS assumptions, which gives the bounds around children's mean potential incomes shown in equation 6.

$$
\begin{gathered}
E[Y^c|D^f < d]Pr(D^f < d) + E[Y^c|D^f = d]Pr(D^f \geq d) \\
\leq E[Y^c(d)] \leq \\
E[Y^c|D^f = d]Pr(D^f \leq d) + E[Y^c|D^f > d]Pr(D^f > d)
\end{gathered} \tag{6}
$$

As indicated in Manski and Pepper (2000), the combined MTR-MTS assumption has a testable implication which is shown in equation 7.

$$
\begin{array}{ccccc}
E[Y^c(d_2)|D^f = d_2] & \overset{MTS}{\geq} E[Y^c(d_2)|D^f = d_1] \overset{MTR}{\geq} & E[Y^c(d_1)|D^f = d_1] & & for\, d_2 > d_1 \\
= & & = & & \\
E[Y^c|D^f = d_2] & & E[Y^c|D^f = d_1] & &
\end{array} \tag{7}
$$

Under the MTS-MTR assumption childen's observed mean income should be weakly increasing in father's realized income category. The top panel in Figure 1 shows that in our data set the mean of children's observed income is strictly increasing in father's income category, consistent with the testable implication in equation 7.[7]

Next we combine the MTR and MTS assumption with the two background characteristics, parents' level of education and childhood neighborhood income, and assume that the MTR and MTS assumptions hold conditional on these two background characteristics. Equation 8 shows the conditional MTR assumption, which states that among children that grew up in a neighborhood with the same income level and have parents with the same level of education, an increase in their father's income will on average not reduce the adult incomes of these children.

$$
E[Y^c(d_2)|D^f = d, N = n, E = e] \geq E[Y^c(d_1)|D^f = d, N = n, E = e] \quad \forall \quad d, n, e, \;\; d_2 > d_1 \tag{8}
$$

It is important to note that the MTR assumption should hold conditional on realized neighborhood income and not conditional on potential neighborhood income. It is possible that part of the effect of increasing father's income on the income of the children is via neighborhood income, i.e. parents moving to a higher income neighborhood due to an increase in their income. This potential channel is not excluded by the conditional MTR assumption. Instead, the MTR assumption should hold within subpopulations of children that have the same *observed* childhood neighborhood income level and parents with the same *observed* level of education.

The conditional MTS assumption is shown in equation 9.

$$
E[Y^c(d)|D^f = d_2, N = n, E = e] \geq E[Y^c(d)|D^f = d_1, N = n, E = e] \quad \forall \quad d, n, e, \;\; d_2 > d_1 \tag{9}
$$

[7] Since children's mean income is clearly strictly increasing in father's income category, there is no need to perform a formal statistical test to show that equation 7 holds in our dataset.

In the hypothetical scenario in which we give all fathers an income within the same category d, the conditional MTS assumption states that within each subpopulation defined by childhood neighborhood income and parents level of education, the mean income of children with initially high income fathers ($D^f = d_2$) is not lower than the mean income of children with initially low income fathers ($D^f = d_1$). For this assumption to hold, any differences in unobserved characteristics between (the offspring of) high and low income fathers within each subgroup should be consistent with the conditional MTS assumption, i.e. children with high income fathers should have weakly better unobservables compared to children of low income fathers.

The conditional MTR-MTS bounds are obtained by first computing the MTR-MTS bounds shown in equation 6 within each subgroup defined by neighborhood income and parent's education and next taking the weighted average over these lower and upper bounds as shown in equation 10.

$$\sum_{n \in N, e \in E} P(N=n, E=e) \cdot MTSMTR\,LB_{E[Y^c(d)|N=n,E=e]}$$

$$\leq E[Y^c(d)] \leq \tag{10}$$

$$\sum_{n \in N, e \in E} P(N=n, E=e) \cdot MTSMTR\,UB_{E[Y^c(d)|N=n,E=e]}$$

Also the conditional MTS-MTR assumption has a testable implication; children's observed mean incomes should be weakly increasing in father's realized income category within each subgroup defined by the level of parent's education and the level of neighborhood income. We have 4 levels of parental education and 10 levels of neighborhood income, which gives in total 40 subgroups. Within each subgroup we tested if children's mean income is weakly increasing in father's income category by performing 11 one-sided difference-in-means tests, resulting in 440 p-values.[8] At a 5 percent significance level we would expect to reject the Null hypothesis while it is true in about 22 tests (0.05*440=22), it is therefore reassuring that the p-value is only smaller than 0.05 in 3 tests. These standard p-values are however not robust to multiple testing. We therefore also computed Romano and Wolf (2005) step-down adjusted p-values which are robust to multiple hypothesis testing. All 440 Romano-Wolf p-values are larger than 0.99, which implies that we do not reject the conditional MTS-MTR assumption in our data.

As a final step, we combine the conditional MTR-MTS assumption with using neighborhood income and parents' level of education as MIV's. We combine the two MIV's by using the following double-MIV assumption

$$E[Y^c(d)|N=n_2, E=e_2] \geq E[Y^c(d)|N=n_1, E=e_1] \quad \forall \quad d,\ n_2 \geq n_1,\ e_2 \geq e_1 \tag{11}$$

Under this assumption, the potential incomes of the children with high educated parents that

[8]Father's income is divided into 12 categories, and we performed tests with the Null hypothesis that $E\left[Y^c|D^f=d,N=n,E=e\right] = E\left[Y^c|D^f=d+1,N=n,E=e\right]$ versus the alternative that $E\left[Y^c|D^f=d,N=n,E=e\right] > E\left[Y^c|D^f=d+1,N=n,E=e\right]$ for $d=1,..,11$.

grew up in a high income neighborhood ($N=n_2, E=e_2$) are on average not lower than the mean potential income of children with lower educated parents ($N=n_2, E=e_1$), from lower income neighborhoods ($N=n_1, E=e_2$) or both ($N=n_1, E=e_1$). Note that the double MIV assumption does not sign the difference in mean potential outcomes of children with high educated parents from a low income neighborhood ($N=n_1, E=e_2$) and the mean potential outcomes of children with low educated parents from a high income neighborhood ($N=n_2, E=e_1$).

By exploiting the assumption in equation 11 we can replace the lower bound on $E[Y^c(d)|N=n, E=e]$ by the maximum of the lower bounds on $E[Y^c(d)|N=n^*, E=e^*]$ for $n^* \leq n$ and $e^* \leq e$, and the upper bound on $E[Y^c(d)|N=n, E=e]$ can be replaced by the minimum of the upper bounds on $E[Y^c(d)|N=n^*, E=e^*]$ for $n^* \geq n$ and $e^* \geq e$. Equation 12 shows the resulting aggregate bounds on $E[Y^c(d)]$, where the lower and upper bounds on the mean potential outcomes within categories defined by the two monotone instruments, are based on the conditional MTS-MTR assumption.

$$\sum_{n \in N, e \in E} P(N=n, E=e)\left[max_{(n^* \leq n, e^* \leq e)} LB_{E[Y^c(d)|N=n^*, E=e^*]}\right]$$

$$\leq E\left[Y^c(d)\right] \leq \tag{12}$$

$$\sum_{n \in N, e \in E} P(N=n, E=e)\left[min_{(n^* \geq n, e^* \geq e)} UB_{E[Y^c(d)|N=n^*, E=e^*]}\right]$$

3.6 Estimating bounds on the average causal effect

The previous subsection showed how we combine the different assumptions to obtain bounds around children's mean potential incomes, $E\left[Y^c(d)\right]$. Equation 13 shows how to use these resulting bounds to construct bounds around the average causal effect of an increase in father's income on child income.

$$LB_{E[Y^c(d_2)]} - UB_{E[Y^c(d_1)]} \leq \left(E\left[Y^c(d_2)\right] - E\left[Y^c(d_1)\right]\right) \leq UB_{E[Y^c(d_2)]} - LB_{E[Y^c(d_1)]} \tag{13}$$

We estimate the bounds by replacing the population means and probabilities by their sample counterparts and next construct 95 percent confidence intervals by applying the methods from Imbens and Manski (2004) using bootstrapped standard errors based on 999 replications. As indicated by Manski and Pepper (2009) the estimated bounds based on the MIV assumptions might suffer from finite sample bias, we therefore apply the bootstrap bias-correction method suggested by Kreider and Pepper (2007).[9]

[9] Kreider and Pepper (2007) suggest to estimate the finite sample bias as $\hat{bias} = \left(\frac{1}{K}\sum_{k=1}^{K}\theta_k\right) - \widehat{\theta}$, where $\hat{\theta}$ is the initial estimate of the upper or lower bound, and θ_k is the estimate of the k^{th} bootstrap replication. The bias-corrected MIV-bounds are subsequently obtained by subtracting the estimated biases from the estimated upper and lower bounds.

4 Results

4.1 Main results

We begin by presenting the bounds around children's mean potential incomes, considering various combinations of the assumptions discussed in Section 3. Panel A in Figure 4 illustrates the upper and lower bounds derived from the combination of the MTS and MTR assumptions, represented by green diamonds connected by dashed lines. Additionally, the graph displays the observed mean income of children, indicated by red dots, corresponding to each category of paternal income. The MTS-MTR bounds on children's mean potential incomes are rather wide, encompassing the observed means for all income categories.

This changes when we impose the MTR and MTS assumptions conditional on the two background characteristics, parent's level of education and childhood neighborhood income. Panel B in Figure 4 shows that the conditional MTS-MTR bounds are substantially tighter and indicate a much smaller slope than the one displayed by the association, implying that the observed association is not driven by father's income alone. For the lower income categories the observed mean incomes are smaller than the lower bound on children's mean potential income, while for the highest income categories the observed mean incomes are much bigger than the upper bounds on children's mean potential income.

Finally, Panel C shows that combining the conditional MTS-MTR assumption with using parental education and childhood neighborhood income as monotone instruments tightens the lower bounds at the lower end and the upper bounds at the upper end of father's income distribution. These conditional MTS-MTR-2MIV bounds are repeated in Figure 5 along with 95 percent confidence intervals around the mean potential incomes of children as well as around the observed means. This graph shows that the bounds are precisely estimated and that most of the observed means also fall outside the 95 percent confidence intervals on the mean potential incomes of children.

As explained in Section 3.6 the bounds on the mean potential incomes can be used to obtain bounds on the average causal effect of an increase in father's income on child income. Row (i), column (3) in Table 1 shows that increasing father's income from the bottom five percent to the top five percent of the income distribution is associated with a rise in child income by on average 419,931 NOK ($52,239). However, the estimated bounds in column (5) show that this increase in paternal income yields an average causal effect ranging between 3,396 NOK ($422) and 212,464 NOK ($26,431). Although the lower bound is statistically significantly different from zero, the upper bound is notably smaller than the intergenerational association. By dividing the upper and lower bounds on the average causal effect by the observed difference in means, it is revealed that the causal component explains at least 1 percent and at most 51 percent of the corresponding intergenerational association in incomes.

Figure 4. Bounds around children's mean potential incomes, under various assumptions

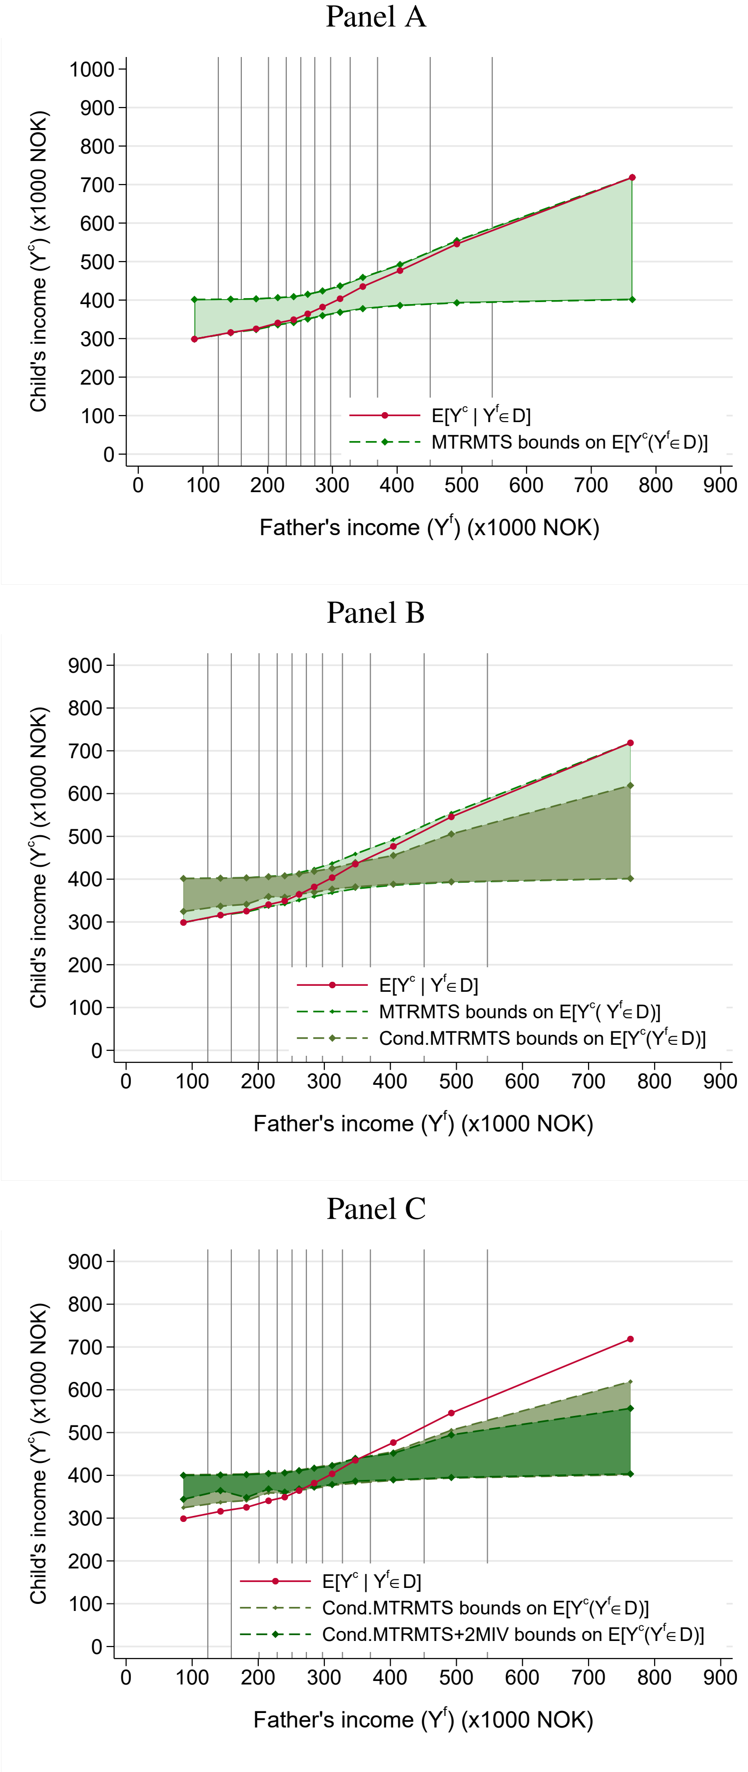

Note: Child's and father's income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds based on the MIV assumption are bias-corrected using the method proposed by Kreider and Pepper (2007). Number of observations equals 427,512.

Table 1. Cond.MTS-MTR-2MIV bounds on the average causal effect of father's income on child income

Treatment	$\triangle Y^f$	$\%\triangle Y^f$	**Diff. in means** ($\triangle Y^c$)	$\%\triangle Y^c$	**ATE** LB	**ATE** UB	**%ATE** LB	**%ATE** UB
	(1)	(2)	(3)	(4)	(5)		(6)	
(i) bottom 5% $\rightarrow$ top 5%	676.68	779.67	419.93	140.59	3.40	212.46	0.85	61.36
			(384.28 455.58)	(128.34 152.84)	(0.71	244.51)	(0.17	72.34)
(ii) bottom 5% $\rightarrow$ 5th decile	175.07	201.72	65.65	21.98	0	66.92	0	19.24
			(60.50 70.81)	(20.04 23.92)	(0	88.33)	(0	25.98)
(iii) 5th decile $\rightarrow$ top 5%	501.60	191.55	354.28	97.24	0	188.28	0	51.03
			(318.54 390.02)	(87.26 107.21)	(0	213.13)	(0	57.80)

Note: Child's and father's income are measured in 1000 NOK deflated to 2015. Bounds on ATE are based on the combination of the conditional MTS-MTR assumption and using parental education and childhood neighborhood income as MIV's (Cond. MTS-MTR-2MIV bounds). Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals in parentheses are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 427,512.

Figure 5. CondMTS-MTR-2MIV bounds around children's mean potential incomes

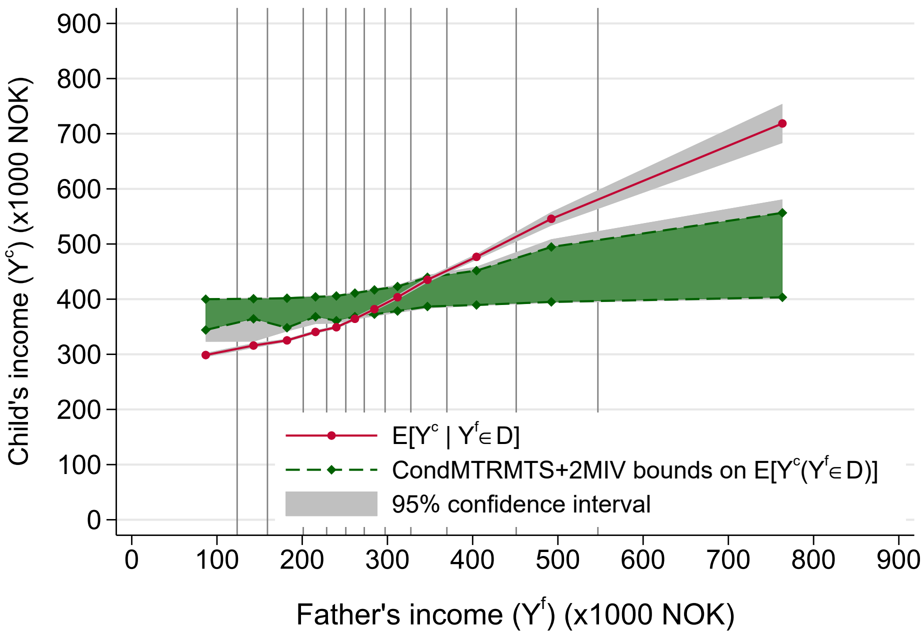

Note: Child's and father's income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 427,512.

Many prior studies estimating intergenerational income mobility concentrate on estimating intergenerational income elasticities rather than changes in income levels. Based on the observed difference in means and estimated bounds on the ATE depicted in row (i) of Table 1, we can also compute approximate intergenerational elasticities. The average percentage increase in father's income associated with increasing father's income from the bottom 5 percent to the top 5 percent equals 779.67 percent. This increase in father's income is associated with an increase in child income by 140.59 percent, which gives an intergenerational income elasticity of 0.18. If we instead use the bounds on average causal effect, the average percentage increase in child income is at most 61.36 percent,[10] which gives an upper bound on the causal intergenerational income elasticity of 0.08.

Figure 5 further illustrates that the slope of the intergenerational association is flatter at the lower end of the father's income distribution, indicating that the relation between father's income and the income of his offspring is nonlinear. This observation is corroborated by rows (ii) and (iii) in Table 1, which demonstrate that increasing father's income from the bottom 5 percent to the fifth decile is associated with an average increase in child income by 21.98 percent, whereas increasing father's income from the fifth decile to the top 5 percent is associated with an average increase in child income by 97.24 percent. Using the corresponding percentage increases in father's income shown in column (2), this implies intergenerational income elasticities of 0.11

[10]This is computed as follows: $100 * \left(UB_{E[Y^c(D^f=top5\%)]} - LB_{E[Y^c(D^f=bottom5\%)]}\right) / LB_{E[Y^c(D^f=bottom5\%)]}$

and 0.51 respectively. Furthermore, the upper bounds on the average causal effects, displayed in column (5), also indicate that the effect of increasing father's income is smaller at the lower end of the income distribution. However, as the bounds on the ATE's overlap, it is not possible to draw strong conclusions regarding differences in the causal effect of father's income along the income distribution.

4.2 Robustness checks

The main results in Figure 5 and Table 1 show that the causal effect of father's income on child income is substantially smaller than suggested by the intergenerational association. In the current subsection we investigate whether these main results are sensitive to the number of categories of the neighborhood income MIV, and to using after-tax income instead of taxable income.

Categories of the neighborhood income MIV As explained in Section 3.3, we categorize neighborhood income into 10 groups; eight groups each spanning 10,000 NOK, with an additional top and bottom category. Increasing the number of categories could potentially lead to tighter bounds. However, this could also result in broader confidence intervals due to fewer observations within each category. Subdividing into too many categories might even lead to empty cells, making it impossible to estimate the bounds.[11]

In this subsection, we re-evaluate our primary analysis by altering the number of categories of the neighborhood income MIV. We start with creating 10 categories, but in contrast to our main analysis we now use 10 deciles of neighborhood income, instead of categories spanning 10,000 NOK. Next, we form six categories, four of which span 20,000 NOK, plus a top and bottom category. We then further divide neighborhood income into 18 categories—14 spanning 5,000 NOK, alongside separate categories for the highest and lowest neighborhood income levels. These categorizations are illustrated by the grey vertical lines in the histograms in Figure A.1 in the appendix.

Figure 6 displays the conditional MTS-MTR-2MIV bounds on the mean potential outcomes using the three alternative categorizations. The bounds we obtain using deciles of neighborhood income are slightly wider, but otherwise very similar to the main results in Figure 5. The bounds based on six categories of neighborhood income are a bit wider compared to our primary results using 10 categories, but also here the difference is minimal.

The bounds derived from 18 categories are noticeably tighter, though they are less precise, particularly at the lower end of father's income distribution. These latter findings suggest that increasing father's income from the bottom 5 percent to the top 5 percent of the income distribution raises child income on average by at least 8,823 NOK ($1,098) and at most 172,848

[11] As indicated in (Manski and Pepper, 2000, 2009), if the number of observations within the sub-groups defined by the MIV's becomes too small, this could lead to finite sample bias. In our main analysis the number of observations in the 40 categories defined by the two MIV's ranges between 1,843 and 31,161, which implies that finite sample bias is unlikely an important issue. We nevertheless apply the bootstrap bias-correction method suggested by Kreider and Pepper (2007) to correct for possible finite sample bias.

NOK ($21,502), equivalent to an increase by at least 2 and at most 49 percent. Based on these results, the causal component accounts for between 2 and 41 percent of the observed difference in means, and the causal intergenerational income elasticity is at least 0.003 and at most 0.06. Despite achieving more informative bounds with 18 categories of neighborhood income, we use 10 categories in our main analysis, because using more categories results in subgroups with too few observations in the analysis by gender which we discuss in Section 4.3.[12]

After tax income Consistent with most research on intergenerational income mobility, our primary analysis uses pre-tax income. However, as Landersø and Heckman (2017) have noted, estimates of the intergenerational transmission of income may be sensitive to the income measure used. Within the Norwegian tax records, we have access not only to taxable income but also to deductions and taxes, which enables us to calculate after-tax income. Due to the wealth tax, a small subset of individuals with high wealth might have positive pre-tax but negative after-tax income. We exclude these individuals from our analysis sample, which decreases the total number of observations by 922 (0.2 percent). Typically, after-tax income is about 68 percent of taxable income.

Figure 7 displays the conditional MTS-MTR-2MIV bounds using after-tax income for both fathers and children. Although the axis scales differ from those in Figure 5, the graphical representations are rather similar. When considering after-tax incomes, the estimated average causal effect of increasing a father's income from the bottom five percent to the top five percent of the income distribution on a child's after-tax income equals at least 2,138 NOK ($266) and at most 96,435 NOK ($11,997). This represents at least 1 and at most 48 percent of the observed intergenerational association in after-tax incomes, which is thus very similar to what we find for pre-tax incomes.

4.3 Gender Heterogeneity

The analyses discussed in the previous subsections combined sons and daughters into a single group. While most previous research on intergenerational income mobility focuses exclusively on fathers and sons, there is an emerging body of literature that examines gender differences in (trends of) intergenerational income persistence. Raaum et al. (2008) estimate intergenerational earnings elasticities for Norway, Denmark, Finland, the UK, and the US, finding that in all these countries, the elasticity of individual earnings with respect to parents' earnings is higher for sons than for daughters. More recent studies (Markussen and Røed, 2020; Ahrsjö et al., 2023; Brandén et al., 2023) analyze trends in intergenerational rank correlations. These studies find that while intergenerational rank persistence (based on parents' or father's earnings rank) is greater for sons than for daughters, this difference is diminishing over time. Conversely, Davis

[12] When we perform the analysis by gender using 18 categories of neighborhood income, the number of observations in some of the categories become too small, resulting in empty categories in part of the bootstrap replications.

Figure 6. Cond.MTS-MTR-2MIV bounds using different categories of neighborhood income

10 deciles

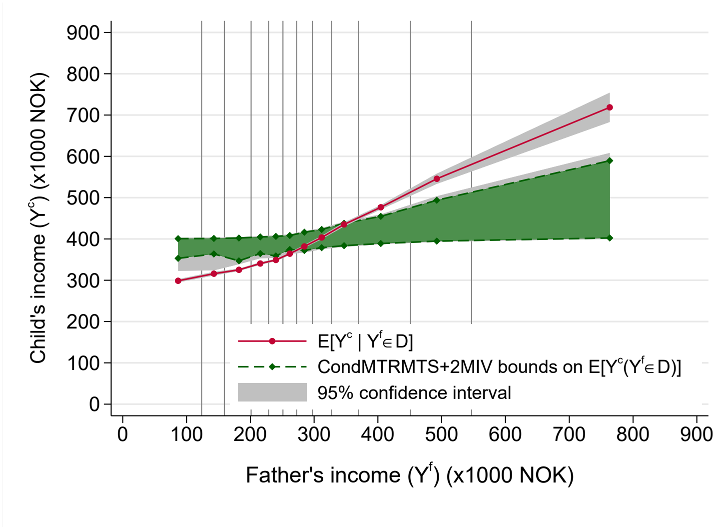

6 categories

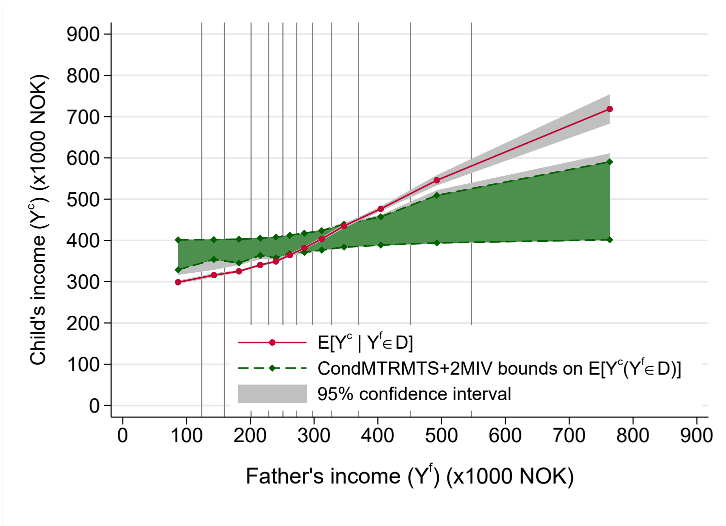

18 categories

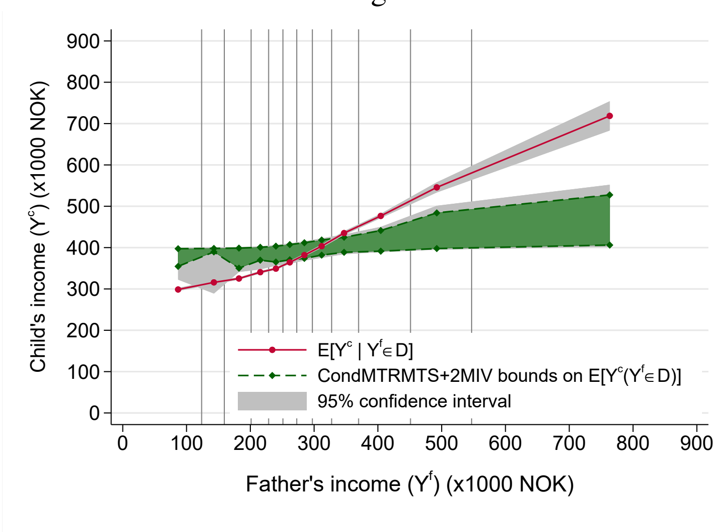

Note: Child's and father's income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 427,512.

Figure 7. Cond. MTS-MTR-2MIV bounds, using after-tax incomes

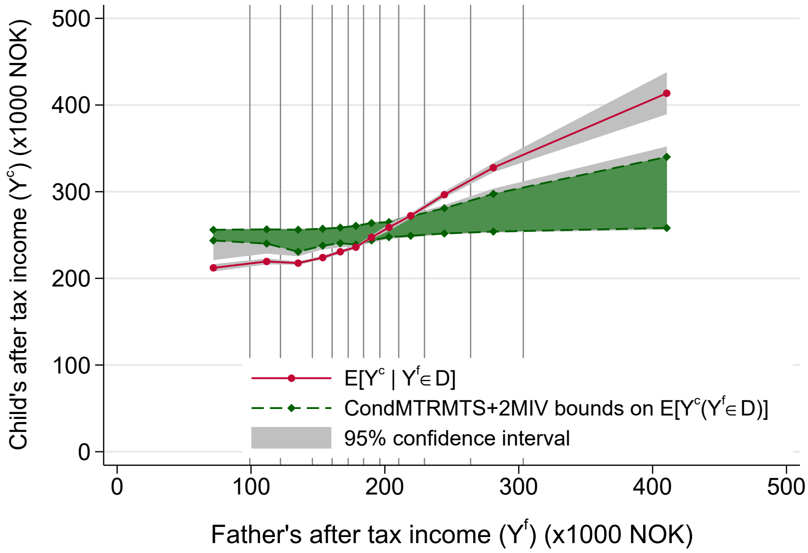

Note: Child's and father's after tax income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 426,590.

and Mazumder (2024), using data from the NLS in the U.S., find that the rank-rank slope based on individual income was initially higher for daughters. However, due to a strong upward trend for sons and a slight downward trend for daughters, the intergenerational rank correlation is higher for sons in later cohorts. Although the literature on gender differences in (trends in) intergenerational income correlations is expanding, there are, to our knowledge, no previous studies that have estimated gender differences in the causal effect of paternal income on long term child income.

To examine potential heterogeneity in the average causal effect of father's income by gender, we divide the estimation sample and calculate bounds on the mean potential outcomes for sons and daughters separately. The results are shown in Figure 8. The first thing to note is that the observed mean income of sons is consistently higher than that of daughters across the entire distribution of father's income. Secondly, we find that the increase in sons' income associated with an increase in father's income is substantially larger than the corresponding increase for daughters. Table 2 shows that increasing a father's income from the bottom 5 percent to the top 5 percent of the income distribution is associated with an average increase in sons' income by 593,743 NOK ($73,862). In contrast, the same increase in father's income is associated with an average increase in daughters' income by 241,722 NOK ($30,070). Using the corresponding increases in percentages shown in column (2) of Table 2, gives an intergenerational income elasticity of 0.217 for sons and only 0.126 for daughters.

Figure 8. Gender heterogeneity: individual income

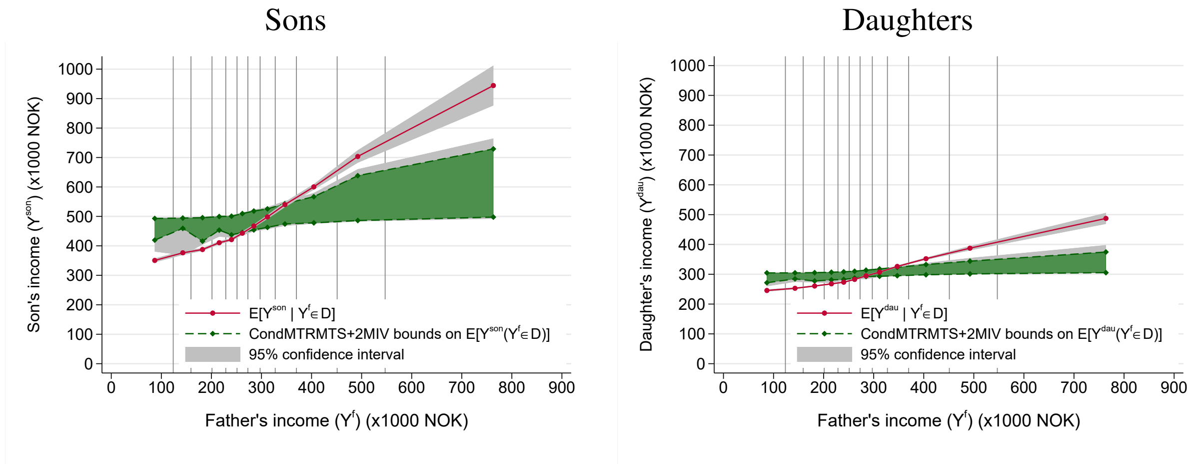

Note: Child's and father's income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 216,819 for sons and 210,693 for daughters.

Table 2. Gender differences in the ATE of increasing father's income from bottom 5% to top 5%

	Diff. in means $\triangle Y^c$	$\%\triangle Y^c$	**ATE**		**%ATE**	
			LB	UB	LB	UB
	(1)	(2)	(3)		(4)	
Son						
(i) individual income	593.74	169.36	4.63	309.58	0.94	72.71
	(526.06 661.43)	(149.71 189.01)	(0	364.15)	(0	88.20)
(ii) household income	730.18	144.93	4.75	400.16	0.68	67.80
	(661.80 798.55)	(131.04 158.82)	(0	458.73)	(0	79.60)
Daughter						
(iii) individual income	241.72	98.41	0.73	102.96	0.24	37.30
	(222.90 260.54)	(90.40 106.42)	(0	128.50)	(0	46.50)
(iv) household income	579.78	107.82	2.98	312.12	0.42	50.91
	(545.43 614.12)	(100.56 115.07)	(0	371.39)	(0	60.85)

Note: Child's and father's income are measured in 1000 NOK deflated to 2015. Bounds on ATE are based on the combination of the conditional MTS-MTR assumption and using parental education and childhood neighborhood income as MIV's (cond. MTS-MTR-2MIV bounds). Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals in parentheses are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 216,819 for sons and 210,693 for daughters.

Figure 8 further illustrates that the bounds on the mean potential outcomes for daughters are much tighter and flatter than for sons. Table 2 shows that an increase in father's income from the lowest to the highest category has an average causal effect on daughter's income between 733 NOK (\$91) and 102,964 NOK (\$12,809). The corresponding upper bound on the average causal effect for sons equals 309,580 NOK (\$38,512) which is 3 times larger. Although both for sons and daughters the estimated lower bounds on this average causal effect are not significantly different from zero, the upper bounds differ substantially and indicate that father's income has at most a very small effect on the income of his daughter, implying a causal intergenerational income elasticity of at most 0.048.[13]

Figure 9. Gender heterogeneity: The effect of the sum of father's and mother's income

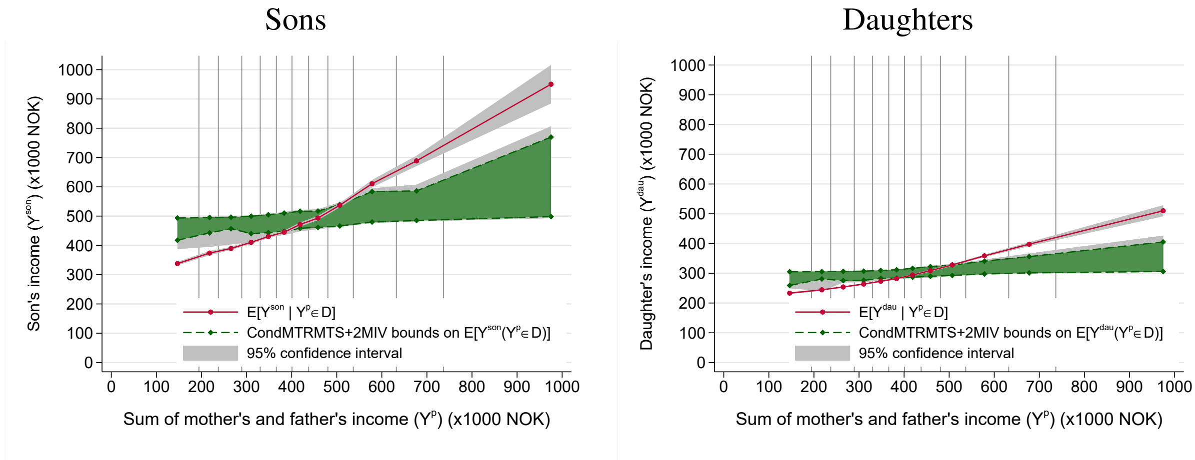

Note: Child's and parents' income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 212,250 for sons and 206,141 for daughters.

One possible reason for the observed differences by gender is that until now we focused on the causal effect of father's income. It might be that fathers are more important for sons while mothers are more important for daughters. Although we also observe long time series of mother's income, using the nonparametric bounding approach to estimate bounds on the average causal effect of mother's income on child income is complicated due to the fact that many mothers do not work or work part time. Especially if (high ability) mothers married to high income fathers decide not to work (or to work less) this can violate the MTS assumption. We therefore do not estimate bounds on the causal effect of mother's income alone, but instead take the sum of father's and mother's income and investigate if the effect of this combined parental income differs between sons and daughters. The results are shown in Figure 9. Apart from the fact that the horizontal axes are longer in Figure 9, the graphs look very similar to the graphs in Figure 8 indicating that the observed gender differences are not caused by the focus on father's income.

Previous studies (Chadwick and Solon 2002; Ermisch et al. 2006; Raaum et al. 2008; Holmlund 2022) have indicated that assortative mating in the child generation can influence intergenerational income persistence. To examine whether the gender differences we observe in the estimated associations and bounds on the causal effects persist when accounting for partner income, we next use the sum of the income of the child and their partner as the outcome variable.[14] Figure 10 and rows (ii) and (iv) in Table 2 presents the results using this new income

[13]The upper bound on the causal intergenerational income elasticity for sons equals 0.093.

[14]For each of the years the child is ages 40 to 50 we take the sum of the taxable income of the child and the

definition.

The red dots connected by red lines in Figure 10 indicate that intergenerational associations are much more similar between sons and daughters when using household income as the outcome variable. Table 2 reveals that for sons, increasing the father's income from the bottom 5 percent to the top 5 percent of the income distribution is associated with an average increase in household income of 730,178 NOK ($90,834). This is higher than the increase in son's own income (row (i)) in absolute terms but smaller in percentage terms. For daughters, the average increase in household income associated with raising the father's income from the lowest to the highest category is 579,777 NOK ($72,124). Unlike sons, this increase for daughters is substantially higher compared to the rise in their own income, both in absolute terms and percentage terms. Therefore, although the observed intergenerational association remains lower for daughters, the differences are considerably smaller when the partner's income is included.

Figure 10 and Table 2 also show that the bounds around the mean potential outcomes and the average causal effects are much more similar for sons and daughters when we focus on the sum with partner's income. The average causal effect of an increase in father's income from the bottom 5 percent to the top 5 percent on household income is between 0.68 and 67.80 percent for sons and between 0.42 and 50.91 percent for daughters, implying upper bounds on the causal intergenerational elasticities of 0.087 for sons and 0.065 for daughters. For both genders we find that the causal component is at most 54-55 percent of the intergenerational association between father's income and household incomes.[15] The comparison of the results in Figures 8 and 10 and the results in Table 2 thus indicate that assortative mating plays an important role in the intergenerational transmission of incomes, especially for daughters.

Figure 10. Gender heterogeneity: Sum of child and partner's income

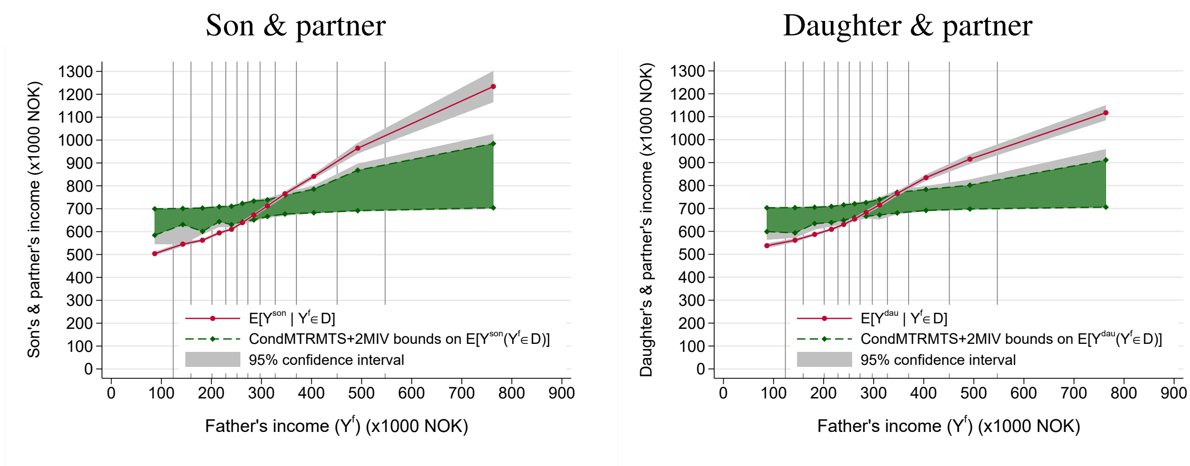

Note: The sum of child and partner's income and father's income are measured in (x1000) NOK deflated to 2015. Estimated upper and lower bounds are bias-corrected using the method proposed by Kreider and Pepper (2007). 95% confidence intervals are based on the method described in Imbens and Manski (2004) and are based on 999 bootstrap replications. Number of observations equals 216,819 for sons and 210,693 for daughters.

taxable income of the child's cohabitant (married and unmarried).

[15] 312,199/579,777=0.54 for daughters and 400,163/730,178=0.55 for sons.

5 Conclusion

Despite Norway's status as an egalitarian country with an extensive social safety net, we still observe a strong intergenerational association between a father's income and the income of his offspring. If this strong association reflects a causal effect, redistributive policy measures can have lasting consequences by affecting not only the current generation but also future generations. However, our estimated bounds on the mean potential outcomes and average causal effects show that at least half of the observed association is due to selection rather than a causal effect.

We also find substantial differences in the intergenerational transmission of incomes between sons and daughters. For sons, increasing the father's income from the bottom 5 percent to the top 5 percent of the income distribution is associated with an average increase in son's income of 169 percent, implying an observed intergenerational income elasticity of 0.217. However, the upper bound on the average causal effect of this increase in father's income is only 73 percent, suggesting a causal intergenerational elasticity of at most 0.093. Compared to previous studies that decomposed the intergenerational income elasticity, this upper bound is informative and lower than the elasticity between the incomes of adoptive fathers and adoptive sons found by Björklund et al. (2006) (0.17), the twin fixed effect estimate by Amin et al. (2011) (0.12), and the upper bound on the causal intergenerational income elasticity found by Lefgren et al. (2012) (0.11).

For daughters, we find even lower upper bounds on the average causal effects. Increasing the father's income from the bottom 5 percent to the top 5 percent of the income distribution increases daughter's income by at most 37 percent, implying an upper bound on the causal intergenerational income elasticity of 0.048. However, when we change the outcome variable to account for the partner's income, the difference between sons and daughters becomes much less pronounced. Both the association between father's income and the child's income, as well as the bounds on the mean potential outcomes and average causal effects, are much more similar between sons and daughters when using household income as the outcome variable. This indicates that assortative mating in the child generation is an important factor in intergenerational income persistence.To obtain a more comprehensive understanding of the (causal) transmission of incomes from parents to children, it is therefore important to consider the role of the partner and to take into account both individual and household income.

References

Ahrsjö, U., Karadakic, R., and Rasmussen, J. K. (2023). Intergenerational mobility trends and the changing role of female labor. *arXiv preprint arXiv:2302.14440*.

Aizer, A., Eli, S., Ferrie, J., and Lleras-Muney, A. (2016). The long-run impact of cash transfers to poor families. *American Economic Review*, 106(4):935–71.

Amin, V., Lundborg, P., and Rooth, D.-O. (2011). Following in your father's footsteps: A note on the intergenerational transmission of income between twin fathers and their sons.

Becker, G. S. and Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of labor economics*, 4(3, Part 2):S1–S39.

Björklund, A., Lindahl, M., and Plug, E. (2006). The origins of intergenerational associations: Lessons from swedish adoption data. *The Quarterly Journal of Economics*, 121(3):999–1028.

Böhlmark, A. and Lindquist, M. J. (2006). Life-cycle variations in the association between current and lifetime income: Replication and extension for sweden. *Journal of Labor Economics*, 24(4):879–896.

Brandén, G., Nybom, M., and Vosters, K. (2023). *Like Mother, Like Child? The Rise of Women's Intergenerational Income Persistence in Sweden and the United States*. IZA-Institute of Labor Economics.

Chadwick, L. and Solon, G. (2002). Intergenerational income mobility among daughters. *American Economic Review*, 92(1):335–344.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228.

Chetty, R., Hendren, N., and Katz, L. F. (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.

Cooper, K. and Stewart, K. (2021). Does household income affect children's outcomes? a systematic review of the evidence. *Child Indicators Research*, 14(3):981–1005.

Davis, J. M. and Mazumder, B. (2024). The decline in intergenerational mobility after 1980. *Review of Economics and Statistics*, pages 1–47.

Durlauf, S. N. (2004). Neighborhood effects. *Handbook of regional and urban economics*, 4:2173–2242.

Ermisch, J., Francesconi, M., and Siedler, T. (2006). Intergenerational mobility and marital sorting. *The Economic Journal*, 116(513):659–679.

Guryan, J., Hurst, E., and Kearney, M. (2008). Parental education and parental time with children. *Journal of Economic perspectives*, 22(3):23–46.

Haider, S. and Solon, G. (2006). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review*, 96(4):1308–1320.

Holmlund, H. (2022). How much does marital sorting contribute to intergenerational socioeconomic persistence? *Journal of Human Resources*, 57(2):372–399.

Holmlund, H., Lindahl, M., and Plug, E. (2011). The causal effect of parents' schooling on children's schooling: A comparison of estimation methods. *Journal of economic literature*, 49(3):615–51.

Imbens, G. W. and Manski, C. F. (2004). Confidence intervals for partially identified parameters. *Econometrica*, 72(6):1845–1857.

Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119.

Kreider, B. and Pepper, J. V. (2007). Disability and employment: Reevaluating the evidence in light of reporting errors. *Journal of the American Statistical Association*, 102(478):432–441.

Landersø, R. and Heckman, J. J. (2017). The scandinavian fantasy: The sources of intergenerational mobility in denmark and the us. *The Scandinavian journal of economics*, 119(1):178–230.

Lefgren, L., Sims, D., and Lindquist, M. J. (2012). Rich dad, smart dad: Decomposing the intergenerational transmission of income. *Journal of Political Economy*, 120(2):268–303.

Ludwig, J., Duncan, G. J., Gennetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R., and Sanbonmatsu, L. (2013). Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. *American economic review*, 103(3):226–31.

Manski, C. F. (1997). Monotone treatment response. *Econometrica: Journal of the Econometric Society*, pages 1311–1334.

Manski, C. F. and Pepper, J. V. (2000). Monotone instrumental variables: With an application to the returns to schooling. *Econometrica*, 68(4):997–1010.

Manski, C. F. and Pepper, J. V. (2009). More on monotone instrumental variables. *The Econometrics Journal*, 12:S200–S216.

Markussen, S. and Røed, K. (2020). Economic mobility under pressure. *Journal of the European Economic Association*, 18(4):1844–1885.

Mogstad, M. and Torsvik, G. (2023). Family background, neighborhoods, and intergenerational mobility. *Handbook of the Economics of the Family*, 1(1):327–387.

Nybom, M. (2024). Intergenerational income mobility. *Research Handbook on Intergenerational Inequality*, pages 56–72.

Nybom, M. and Stuhler, J. (2016). Heterogeneous income profiles and lifecycle bias in intergenerational mobility estimation. *Journal of Human Resources*, 51(1):239–268.

Raaum, O., Bratsberg, B., Røed, K., Österbacka, E., Eriksson, T., Jäntti, M., and Naylor, R. A. (2008). Marital sorting, household labor supply, and intergenerational earnings mobility across countries. *The BE Journal of Economic Analysis & Policy*, 7(2).

Romano, J. P. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association*, 100(469):94–108.

Sacerdote, B. (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.

A Appendix

Figure A.1. Histograms of neighborhood income with different categorizations

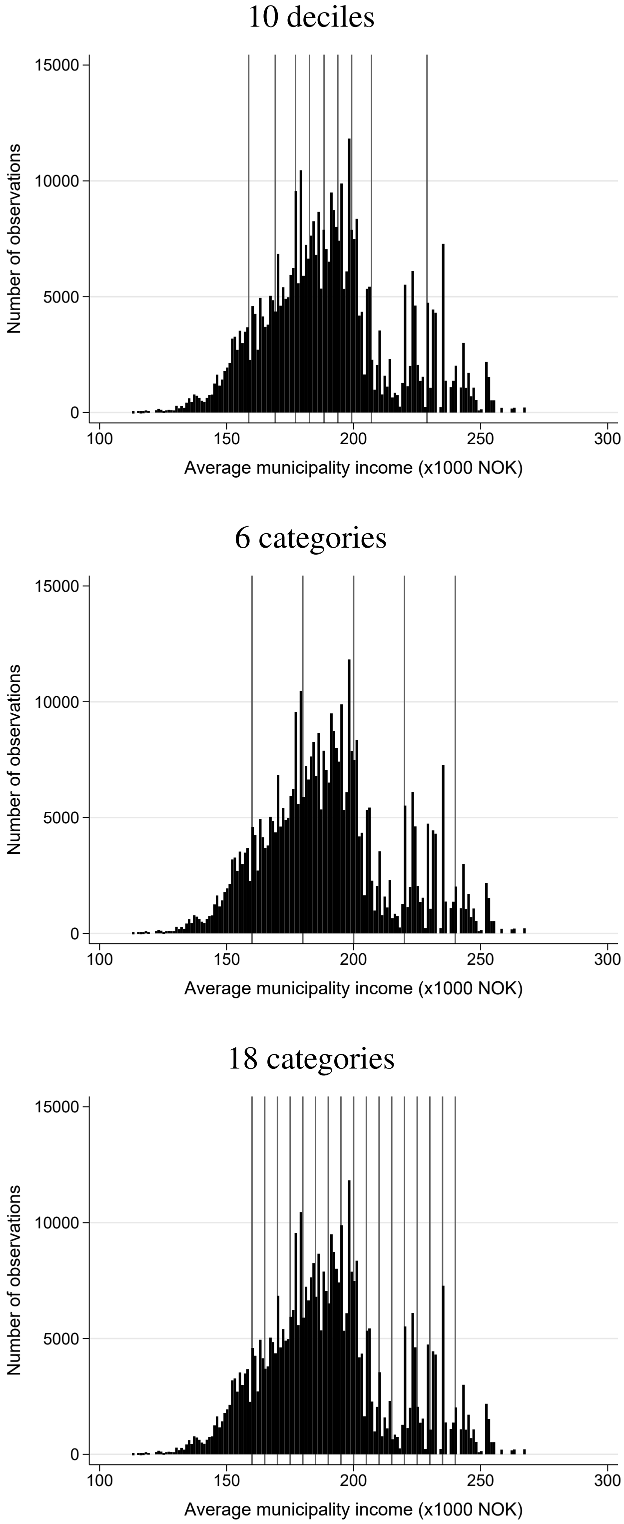

Note: The black bars show the histogram of mean taxable income among prime-aged individuals in the municipality where the child lived at age 16, with the number of observations in each 1000 NOK bin. We take a 3-year average to get a more precise measure and incomes are adjusted to 2015 NOK to account for inflation. The gray vertical lines show how we have defined the neighborhood income MIV. Total number of observations equals 427 512.