Family Background, Neighborhoods and Intergenerational Mobility

Magne Mogstad and Gaute Torsvik
JUNE 2021

FAMILY BACKGROUND, NEIGHBORHOODS AND INTERGENERATIONAL MOBILITY

Magne Mogstad
Gaute Torsvik

May 2021

Family Background, Neighborhoods and Intergenerational Mobility
Magne Mogstad and Gaute Torsvik
May 2021
JEL No. D1,J13,J24,J62

<u>ABSTRACT</u>

This paper reviews the literature on intergenerational mobility. While our review is centered around the large empirical literature on this topic, we also give a brief discussion of some of the relevant theory. We consider three strands of the empirical literature. First, we discuss how to measure intergenerational persistence in various socio-economic outcomes. We discuss both measurement challenges and some notable findings. We then turn to quantifying the importance of family environment and genetic factors for children's outcomes. We describe the pros and cons of various approaches as well as key findings. The third strand is concerned with drawing causal inferences about how children's outcomes are affected by specific features of their family environment. We discuss a wide range of environmental features, including the neighborhoods in which children grow up. We critically assess what conclusions one may and may not draw from certain celebrated studies of neighborhoods and intergenerational mobility.

Magne Mogstad
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
magne.mogstad@gmail.com

Gaute Torsvik
University of Oslo
Oslo
Norway
gaute.torsvik@econ.uio.no

1 Introduction

A large and growing body of evidence has helped policymakers better understand the key drivers of labor market inequality. One of these drivers is the wage premium associated with higher education and, more broadly, with a worker's abilities and skills. Recent research using employer-employee data from many countries documents that a majority of the observed earnings inequality can be attributed to permanent or at least persistent characteristics of individual workers, not the types of firms or industries in which they work (see e.g. Bonhomme et al. (2020)). This research raises several important questions: What are the origins of the large differences in individuals' marketable characteristics and skills? How much does the family that children are born into and the neighborhood they grow up in matter for their economic outcomes?

Empirically it is well documented that economic prosperity tends to persist across generations. Children born to parents with high education or income can expect to do better than children born into poorer conditions. While the degree of intergenerational persistence varies across countries and outcomes, family background is universally a strong predictor of children's outcomes. There is also an emerging perception, and concern, that widening economic inequality is accompanied by an increasing persistence of inequality across generations (OECD, 2018).

There are several reasons why policy makers and researchers are interested in, and concerned about, lower intergenerational mobility. One is fairness. High persistence in socioeconomic status across generations could be interpreted as a sign that economic inequalities within a generation are unfair, caused by disproportionate family resources and unequal opportunities, rather than differences in individual diligence and effort. Another concern is that high intergenerational persistence in economic status is a sign of inefficient use of human resources. If children who are born into difficult circumstances do not get the opportunity to develop their productive potentials, many talents will remain unused or under-developed. A third concern is instability. If children from disadvantaged families perceive that the playing field is tilted in favor of those who are born into riches, there might be less social cohesion and cooperation among citizens.

Although the individual is the key decision maker in economic analysis, economists have long been aware that the preferences, beliefs and constraints of individuals are partly shaped by their family background. Knight (1935) considered the family to be the key economic unit, precisely because it endows individuals with skills and attitudes that matter for the opportunities they have, and choices they make, later in life. Knight also realized that families differ

in their capacity to give their offspring the resources they need to succeed, and hence that family background gave some people important advantages in a competitive labor market.

A few decades later, Becker (1981) developed his treatise of the family as an economic unit and also explored how family endowments and decisions forge intergenerational links between parents' and children's economic outcomes. As better data have become available, there has been a steady growth in empirical studies that measure the strength of intergenerational links in various measures of economic success. To date, the accumulated body of work on intergenerational mobility consists of a relatively small theoretical literature and a large, often quite atheoretical, empirical literature.

We think it is unfortunate that much of the empirical work appears as a catalog of undigested mobility estimates, often unrelated to economic policy and with little if any link to economic theory. A theoretical connection would make it easier to understand and compare the empirical results, such as why family background appears to be more salient in certain places. A firmer grip on why family matters for children's outcomes will also facilitate our understanding of how different policies affect intergenerational mobility, and will make it easier to prescribe policies that help individuals realize their potential in the labor market.

In a hope to move (empirically grounded) theory a bit higher on the research agenda on intergenerational mobility, we start this review article with a discussion of why and how parents may influence the skill and abilities of their offspring. Section 2 presents a theoretical backdrop for how family background may shape skill and abilities that are valued in the labor market, and therefore help explain the inequalities we observe in the labor market.

Next, we turn to empirical studies on intergenerational mobility. This literature is vast and we do not attempt to give a full account of it. We do not cover studies on so-called absolute intergenerational mobility, which is concerned with the likelihood that children have better life outcomes than their parents. We focus instead on the measurement of so-called relative mobility, that is, on how the expected outcomes for children depend on the socioeconomic status of their parents.[1] There is also a large sociological literature on social mobility, which focuses on class and occupational persistence, that we ignore in this review. Even within the topics we cover, our review is more eclectic than comprehensive. Our primary goal is to give an overview and critical assessment of a set of studies that are important or at least influential.

To be concrete, we consider three strands of the empirical literature on intergenerational mobility. One of these, discussed in Section 3, is centered around the measurement of intergenerational persistence in various socio-economic outcomes, such as education and earn-

---

[1]For a discussion of absolute versus relative intergenerational mobility and the relationship between them, see Narayan et al. (2018)

ings. In this section, we discuss both measurement challenges and the main empirical findings. We also document how access to large-scale data from administrative registers allows researches to both get more reliable estimates and to address questions that previously could not be studied.

While it is important to document intergenerational persistence across countries and outcomes, the ultimate goal is to reach beneath the surface and understand why family background matters for life outcomes. This is the goal of the second strand of the literature, discussed in Section 4, where we describe both the pros and cons of various approaches and some key findings. In this section, we first discuss how sibling correlations have been used to obtain an omnibus measure of the role family background plays for educational attainments and for determining labor market outcomes. Sibling correlations reflect the influence of both genetics and family environment. Next, we discuss ways in which researchers have tried to separate these factors.

The third strand of the literature is concerned with estimating the causal effects of specific features of the family environment on children's outcome. For example, what is the causal effect of parents' education and income on offspring's educational attainment and income; to what extent will allowing parents disability benefits reduce their children's attachment to the labor market? In Section 5, we review this literature and discuss how the empirical findings can be interpreted through the lens of basic economic models of intergenerational persistence. In Section 6, we dig deeper into the importance of a specific element of family environment that has recently received a lot of attention, namely the neighborhoods where children grow up. We critically assess what conclusions one can and cannot not draw from a few widely cited studies of neighborhoods and intergenerational mobility.

2 Family background and earnings: theoretical considerations

Children are born with different cognitive and non-cognitive capacities to acquire the kind of knowledge, skills and attitudes—the human capital—that the labor market values. The genes that are transmitted from parents to children may constrain what individuals can achieve in the labor market. The environment they grow up in influences to what extent they reach their potentials.

One element of human capital that parents may influence is the investment in formal education. Economists have typically focused on this type of human capital, but there is an increasing awareness that this perspective is too narrow. As we discuss below, evidence shows that many different types of skills, attitudes and traits are valued in the labor market,

and, moreover, that parents matter for the development of both cognitive and non-cognitive skills at different stages of childhood.

This section gives a short review of the human capital approach to intergenerational mobility.[2] We start with how parents influence formal education before we extend the notion of human capital. At the end, we briefly discuss how parents can influence the labor market outcomes of their children through channels other than the formation of human capital.

2.1 Family resources and investment in schooling: The Becker-Tomes model

In two closely related papers, Becker and Tomes develop a model of the transmission of earnings, assets, and consumption from parents to children (Becker and Tomes, 1979, 1986). The model is based on utility maximization by parents who care about the income or welfare of their children. The model has one period of childhood, one period of adulthood, one child per family (no fertility decisions), and a single parent. Parents begin with income Y_t, a combination of earnings and financial transfers they received from their parents. Parents spend on three items: their own consumption C_t, investments in the human capital of their child I_{t+1}, and financial transfers to their child X_{t+1}. Parents transmit ability or endowment A_{t+1} to their children through a stochastic linear autoregressive process. After observing the child's ability, parents invest in education. Education and ability determine the child's human capital and productivity. In adulthood, workers (grown children) supply labor inelastically.

Parents care about their own consumption and the income available for consumption and investment for their children. The parent's optimization problem is

$$\max_{C_t, I_{t+1}} U(C_t, Y_{t+1}) \tag{1}$$

subject to

$$
\begin{aligned}
Y_t &= C_t + X_{t+1} + I_{t+1}, \\
Y_{t+1} &= w_{t+1} f(I_{t+1}, A_{t+1}) + (1 + r_{t+1}) X_{t+1} + E_{t+1},
\end{aligned}
\tag{2}
$$

and the borrowing constraint

$$X_{t+1} \geq 0, \tag{3}$$

where w_{t+1} is the return to human capital in period $t+1$, $f(\cdot)$ is the human capital produc-

---

[2]Parts of this discussion draw on the review article by Mogstad (2017).

tion function, r is the return on financial assets, and E_{t+1} is the idiosyncratic component of children's income (market luck).

There are two potential links between the income of parents and children in this framework. One of these links is that high achieving parents tend to have good genes, and a fraction of these genes are transmitted to the child. The genes are contained in the endowment variable A in equation (2). However, A contains more than the parents' genes, as we discuss in detail later. The other link is that parents with high earnings tend to invest more in their child's education. They invest more for two reasons: (i) there is a complementary between innate ability and the returns to the investment in human capital and (ii) credit constraints may prevent poor households from investing efficiently in their children. These two transmission mechanisms can attenuate the degree of intergenerational mobility in the economy.

In a more recent paper, Becker et al. (2018) develop a model where inequality in one generation may create more inequality in the next generation. This model shuts down the genetic transmission mechanism by assuming that all parents—irrespective of their own human capital and economic status—have children with the same innate ability. There are two key assumptions in their model. One is that parents with more education and income have higher returns on human capital investments in their children. The other assumption is that there are increasing returns to human capital in the labor market. These assumptions can lead to situations where there is no regression towards the mean income across generations.

If a government is included in these models, it could speed up mobility in two ways, either by imposing a progressive tax and redistribution system that reduces the returns to education, or by providing subsidized education that relaxes the role of the credit constraints poor parents may face when they invest in their child's education (Solon, 2004).

2.2   Human capital beyond the Becker-Tomes model

The Becker-Tomes model captures an important transmission channel that can explain intergenerational persistence of economic outcomes. However, a two period model where investment in schooling is the only choice parents make to influence the future of their children is overly simplistic. Recently, particular attention has been devoted to three assumptions of the Becker-Tomes model: i) investments at any stage of childhood are equally effective, ii) earnings depend on a single skill, and iii) parental engagement with the child is in the form of investment in educational goods, analogous to firm investments in capital equipment. An active body of research suggests that these assumptions are at odds with the data, and that the Becker-Tomes model misses key implications of richer models of intergenerational

mobility. Work by Cunha and Heckman (2007), Heckman and Mosso (2014) and Lee and Seshadri (2019) highlights three ingredients of particular significance for measuring and understanding intergenerational mobility: multiplicity of skills, multiple periods of childhood and adulthood, and several forms of investments.

It is well documented that wages depend on a large set of personal traits, abilities and skills in addition to cognitive skills measured by IQ, or human capital measured by formal education; see Bowles et al. (2001) and Borghans et al. (2008) for evidence. Recently, important progress has been made in accounting for measurement error and in trying to establish causal rather than merely predictive effects of worker characteristics on earnings. The evidence points to the importance of including sufficiently broad and nuanced measures of skills in studies of intergenerational mobility. Both cognitive and non-cognitive skills predict adult outcomes, but they have different relative importance in explaining different outcomes. For example, schooling seems to depend more strongly on cognitive skills, whereas earnings are equally predicted by non-cognitive abilities. For intergenerational mobility, it is important that gaps in many of the relevant skills between socioeconomic groups open up at an early age (Cunha et al., 2006).

In models with multiple periods of childhood and adulthood, the timing of income can be important as it interacts with restrictions on credit markets and the technology of skill formation. Parents may not only be restricted from borrowing against the earnings potential of their children (an intergenerational credit constraint, as in (3)), but also prevented from borrowing fully against their own future earnings (an intragenerational credit constraint). The intragenerational credit constraint induces a suboptimal level of investment (and consumption) in each period in which it binds. How harmful this constraint will be depends on the technology of skill formation and the life-cycle profile of parental earnings. Cunha et al. (2010) estimate the elasticity of substitution parameters for inputs at different periods that govern the trade-off of investment between a child's early years and later years. They present evidence of dynamic complementarities in the production of human capital, implying that early investment in children's skill development will have large returns because it raises the payoffs to future investments. As a consequence, if the intragenerational credit constraint is binding during the early periods, late investments will be lower, even if the parent is not constrained in later periods.

The Becker-Tomes model (and many of the extensions of the model) considers only a single child investment good. Recent evidence, however, points to the importance of allowing for multiple forms of investments, and letting the returns to these investments vary over the life-cycle of the child (Cunha and Heckman, 2007). For example, Del Boca et al. (2014)

develop and estimate a model of intergenerational mobility where parents make a number of specific input choices, ranging from various time inputs to child good expenditures, each with a child age-specific productivity. Their empirical results indicate that both parents' time inputs are important for the cognitive development of their children, particularly when the child is young. In contrast, the productivity of monetary investments in children has limited impacts on child quality no matter what the stage of development.

Doepke and Zilibotti (2017) explore how parental involvement and child development may depend on inequalities in the labor market. In their model, parents decide how much time and resources they want to use on child rearing and development. Their parenting style, whether they choose an authoritarian or permissive style, depends on their temperament and the distribution of earnings in the labor market. Parents may interfere and alter the priorities of their children, for example by "forcing" the child to do her homework. Parents will succumb to this authoritarian parenting style if they are prone to paternalism, that is, if it costs them little to go against the choices and welfare of their kids, or if the earnings premium in education is sufficiently high. According to this model, parents will be less involved in shaping the skills of their kids in an egalitarian economy than they will in an economy where the returns to skills are very steep. Hence, if we assume that rich parents are better equipped with resources and temperament to take the authoritarian parenting style, we should, based on this argument, expect more mobility in egalitarian economies.

2.3   *Opening up the black box of children's endowments*

The endowment variable *A* in (models that build on) the Becker-Tomes framework is supposed to be a composite measure of many factors:

> *"The income of children is raised by their "endowment" of genetically determined race, ability, and other characteristics, family reputation and "connections," and knowledge, skills, and goals provided by their family environment. The fortunes of children are linked to their parents not only through investments but also through these endowments acquired from parents (and other family members)."*(Becker and Tomes, 1979, p. 1153)

Even though children's endowments contain several factors that may deliberately be altered by their parents, such as knowledge, connections and ambitions, these factors are taken to be exogenous in Becker and Tomes (1979, 1986), and consequentially play second fiddle in their analysis.

8

One intergenerational link arises if parents' own choices, outcomes and goals are emulated by offspring because parents convey expectations and set standards for their children. Lindbeck et al. (1999) argue that the uptake of welfare benefits is partly determined by monetary incentives and partly by the strength of the social norm that one should live on one's own work, and that the stigma of violating the norm depends on whether peers, including parents, are dependent on welfare benefits. See also Durlauf and Ioannides (2010) and Manski (2000) for further discussion of this type of expectation or preference interaction, and Dahl et al. (2014) for empirical evidence of how disability benefits reduce offspring labor market attachment.

Another possible source of a child-parent link that may create persistence in outcomes across generations is the educational and labor market information and insights children may obtain from their parents. Becker et al. (2018) captures the essence of this idea by assuming that the returns to education for a child are increasing in the education of the parent. They argue that better educated parents have access to information that allows them to make more efficient investments in their children. It is also possible that local interactions with other advantaged children enhances the returns to formal education. In a more specific context, it has been argued that information and complementarities are the reasons why so many children of medical doctors become doctors themselves (Lentz and Laband, 1989).

Becker suggests another reason why there is a strong intergenerational persistence in the medical profession;

> *"Medical schools have been accused, with some justification, I believe of discrimination against minority groups and favoritism towards relatives of AMA members. Perhaps this explains why doctor's sons more frequently seem to follow in their fathers footsteps than do sons of other professions,* (Becker, 1959, p. 217–218)

It is not only the doctor's child that tends to end up in the same profession as their parents. Corak and Piraino (2011) show that as much as 40% of a cohort of young Canadian men had been employed, at some point in time, at an employer where their father had worked. Kramarz and Skans (2014) use employer-employee linked register data from Sweden and find that parents seems to play a crucial role for whether and where young workers get their first jobs. It does not have to be favoritism that creates a link between the workplace of the parent and offspring. It can be information costs or incentives that induce employers to recruit through strong social ties (Dhillon et al., 2021).

Taken together, the work discussed above suggests that opening up the black box of children's endowments is essential to understand why and how intergenerational outcomes are

linked. We are hopeful that important progress can once again be made by combining theory and empirics, adjusting the theories of intergenerational mobility in light of new evidence and then taken those theories to new data sets.

3 Measurement of intergenerational mobility

We now shift our attention to the empirical literature on measuring intergenerational mobility.

3.1 Measurement and data issues

One way to capture the importance of family background is to measure intergenerational persistence in specific achievements or outcomes, such as schooling, earnings or wealth. The canonical measure of intergenerational persistence is the regression coefficient of children's outcomes on parents' outcomes. This statistic captures to what extent socioeconomic status regresses towards the mean outcome over a generation. Galton (1886) used this model to examine the relationship between the height of parents and children. He concluded that human height regressed towards the population mean by a factor of two-third over a generation.

In economics, the typical specification in analyses of intergenerational earnings (or income) persistence is to regress log earnings (or income) of child i, Y_i^{child}, on the log earnings (or income) of his or her parents , Y_i^{parent}:

$$\log Y_i^{child} = \alpha + \beta \log Y_i^{parent} + \varepsilon_i. \tag{4}$$

The coefficient β is equal to the intergenerational elasticity (IGE). The degree of intergenerational mobility can be measured as $(1 - \beta)$. IGE is a scaled version of the intergenerational correlation (IGC). Using lower case for logs, it follows that IGC $= corr(y^{child}, y^{parent})$ and $\beta = (\text{IGC}) \frac{\sigma_{ychild}}{\sigma_{yparent}}$. Both the dependence structure between parents and offspring outcomes and the marginal distributions matter for the magnitude of IGE. This type of regression model is frequently used to estimate the intergenerational persistence of outcomes in addition to earnings and income, such as educational attainment (Björklund and Salvanes, 2011) and wealth (Boserup et al., 2016).

The first empirical studies on intergenerational mobility in economics reported that earnings regressed faster towards the mean than human height. Becker and Tomes (1986) refer to a handful of studies that use earnings data to estimate IGE and conclude that a reasonable assessment of IGE is around 0.2; only around 20% of the economic advantages (or disadvantages) in one generation are transmitted to the next generation. With so little persistence,

economic success is basically won or lost within a generation; in two generations only 4% of the initial advantage is left.[3]

It soon became clear, however, that measurement errors, unrepresentative samples and other data issues contributed to low IGE estimates. If the goal is to measure to what extent economic privileges and well being persist within families over generations, it seems natural to consider how life time achievements or permanent income (or perhaps consumption) are linked across generations. For this purpose, the ideal data set should contain several years of income for both parents and children, preferably measured around the middle of their careers (Mazumder, 2016).[4] These data are not easy to obtain. In fact, many of the data sets that are used to estimate IGE have no information on parental income at all. Parental income is often predicted based on other covariates (Inoue and Solon, 2010). In the data sets that include income information for the parent generation, income is typically measured only for one or a few years and often late in the parent's career so that it can be linked to their children's earnings.

It is well understood that classical measurement errors in the income of the parents will bias the estimated persistence coefficient towards zero. Another problem is that, if parents and children are observed at different stages in their earnings life cycle (children early and parents late), the IGE in earnings may be biased downwards, since there is a tendency that those with high permanent income have a steeper log earnings profile (Haider and Solon, 2006). Finally, since IGE compares log earnings across generations, individuals with zero earnings are typically dropped from the analysis, which may create bias in the IGE estimates. In practice, this problem is particularly severe if the data contain only a few years of observations for parents or offspring (Mazumder, 2016).

To get around the zero earnings issue, Mitnik and Grusky (2020) suggest that instead of estimating the conditional expected log income of children one should estimate the log value of children's expected earnings, conditional on parents earnings. They denote the coefficient of this regression $\text{IGE}_E = \frac{d\left(\log E\left(Y_i|Y_{j(i)}=y\right)\right)}{d\log y}$. However, this measures a different parameter than the IGE, and its justification is unclear, except that it is also defined for those with zero income. It does, for example, not approximate the expected (average) welfare or utility of

---

[3]This simple calculation assumes that intergenerational mobility is a Markov process: there is no direct effect of grandparents on grandchildren; the link goes via parents. There is a small but growing body of work on multigenerational transmission of socioeconomic status (see e.g. Solon (2018)). While this work is very promising, we do not review it in this article.

[4]The relevant measure of income will naturally depend on the degree of income variability and uncertainty individuals face and also on whether they are credit constrained. The results in Carneiro et al. (2021) suggest that the timing of parental income matters for child development.

children conditional on parent earnings if we assume diminishing marginal utility of earnings (for example if utility can be represented by log earnings).

As another alternative to IGE, the rank-rank regression has become increasingly popular, especially after the work of Chetty et al. (2014a). Comparing ranks across generations allows researchers to include children and parents with zero labor market earnings. Another potential advantage is that the rank-rank regression coefficient isolates the dependence structure across generations, essentially by making the marginal distributions uniform and, thus, invariant to changes in inequality within a generation.[5] Whether this is a pro or a con depends on the question of interest. For example, Chetty et al. (2014b) use rank-rank regressions to examine intergenerational mobility in the U.S. over time. They find that children entering the labor market today have the same chances of moving up in the income distribution (relative to their parents) as children born in the 1970s. However, since inequality has increased, family background has a larger impact on expected income today than in the past. In other words, intergenerational persistence has increased if one consider individuals' income instead of their ranks. For the same reason, one may want be cautious in interpreting the results from rank-rank regressions across countries, as moving up in the income distribution may have a very different impact on an individual's income and welfare depending on the cross-sectional inequality in her country.

Data limitations have constrained the empirical knowledge of intergenerational persistence in earnings or income. Educational attainments, or occupational data, are in many ways easier to measure and to link across generations. It is easier to recall the education level of parents and grandparents, or their occupation, than it is to give an assessment of their earnings. Life cycle issues are also less relevant since most individuals have completed their education before they are 30. Studies that compare intergenerational persistence across rich and poor countries (Narayan et al., 2018), or studies that compare changes in intergenerational persistence in socioeconomic outcomes over a long horizon (Modalsli, 2017), therefore tend to focus on occupations or educational attainment. On the other hand, it can be challenging to make meaningful comparisons of education levels across time and especially across countries, because the educational system may change over time, and different countries may be at different stages in this transition process (Karlson and Landersø, 2021).

---

[5]It is, however, not invariant to changes in inequalities that are independent of parental rank. To see this dependence suppose that in society A female workers are paid half of the wages of equally productive males. In society B there is no such discrimination. Since the gender of a child is basically random and independent of parental rank, there will be more intergenerational mobility in the gender discriminating society A than in equal pay society B. See Gandil (2019) for a further discussion.

3.2   Selected findings

Four patterns emerge from the large number of studies that estimate IGEs in earnings (income) and education. First, there are systematic differences in intergenerational mobility across developed countries. Mobility, as measured by the IGE, appears to be relatively low in the US, a bit higher in the UK and in continental Europe, and highest in Canada and the Nordic countries (Corak, 2006; Björklund et al., 2009a; Blanden, 2013). Intergenerational mobility in education attainment follow the same pattern, although the gap between low mobility (the US) and high mobility countries (the Nordics) is smaller for education than it is for income (Landersø and Heckman, 2017; Karlson and Landersø, 2021). A second pattern that emerges from international comparisons is a negative relationship between mobility and income inequality. This relationship has its own name, the Great Gatsby Curve. A third pattern is that intergenerational mobility is considerably lower in the developing world than in developed countries (Narayan et al., 2018). The fourth pattern is that intergenerational mobility varies not only across countries, but also across regions within a country (Stuhler, 2018).

Before discussing some of these patterns in greater detail, it is important to observe that (in part due to the measurement and data issues described above) there is considerable uncertainty associated with the IGE estimates, and especially with their comparison across time and place. This is vividly illustrated by the disarray of mobility estimates in the U.S.

Early evidence from the U.S. indicated an IGE of around 0.2 and, thus, portrayed America as a land of opportunity. However, when Solon (1992) and Zimmerman (1992) included additional years of data on fathers' earnings to reduce the measurement error of permanent income, IGE increased to about 0.4. Mazumder (2005) used 15 years of income data and estimated IGE to be around 0.6 for both men and women. Dahl and DeLeire (2008) use a small sample of administrative income data and find estimates that are very sensitive to how fathers with zero income are treated. Their estimate of earnings IGE vary from 0.26 to 0.63 for men and from 0 to 0.27 for women. Chetty et al. (2014a) use the full population of administrative tax records in the period 1996 to 2012 and estimate the IGE of income be around 0.34 for both male and female children. Mazumder (2016) argues that the low IGE of Chetty et al. (2014a) is partly because they measure children's income only for two years and very early in their career, but also because there is a substantial fraction (7%) that do not file a tax return. The results in Chetty et al. (2014a) are very sensitive to how one treats these non-filers and what choices one makes about the tails of the income distribution. For example, if one estimates the IGE on a sample that excludes the bottom and top 10%, the IGE

13

becomes as large as 0.45. Landersø and Heckman (2017) also estimate the IGE in the US to be 0.45 for income (and 0.29 for earnings). Mitnik et al. (2015) use a random sample of tax data that gives a longer time series on income. This enables the authors to measure children's earnings over several years and later in their career than Chetty et al. (2014a). Mitnik et al. (2015) then find an IGE of just above 0.5 for men and just below 0.5 for women.

USA, the Nordics and the Great Gatsby Curve.   Given the wide rang of IGE estimates within one country it may seem overly optimistic to think we can learn much from making cross country comparisons of intergenerational mobility. Nevertheless, such comparisons are common. The most convincing comparative studies standardize data sets and methodologies to estimate and compare IGE in different countries. Many of the early studies compared intergenerational transmissions of income in egalitarian Nordic welfare states with the more laissez-faire oriented economies in the U.S. and the U.K. Björklund and Jäntti (1997) is an early example of such studies, comparing intergenerational mobility in Sweden and the U.S. Their motivation was to examine the widespread argument that the *intra*generational economic inequality in the U.S. was accompanied by higher *inter*generational mobility. The results in Björklund and Jäntti (1997) did not support the idea of the U.S. as a country with high mobility. They estimated a lower IGE in the US than in Sweden. However, due to small samples the estimates were imprecise and the authors could not draw any firm conclusions.

Armed with larger data sets, Jantti et al. (2006) and Bratsberg et al. (2007) compared intergenerational mobility in the United States, the United Kingdom, Denmark, Finland, Norway and Sweden. Using earnings data for fathers and sons, they find that intergenerational persistence is highest in the United States (IGE=0.52) and lowest in Denmark (IGE=0.07). Based on these studies, higher intragenerational inequality is associated with *lower* intergenerational mobility.

Corak (2006) includes data from more countries and confirms a strong negative relationship between cross-sectional inequality and intergenerational mobility. This relationship was denoted the Great Gatsby Curve by Krueger (2012). Figure 1 uses more recent data and plots this relationship for a larger group of countries. The figure depicts a wide variation in IGEs, from low 0.12 (Denmark) to a high 0.76 (Colombia). The negative association between intergenerational mobility and inequality is not very clear. In fact, one interpretation of the data is that it contains three clusters of countries. If we remove the emerging economies with relatively high inequality and the Nordic countries with relatively low inequality, and focus on the majority of the OECD countries in the middle, one could argue that higher mobility is associated with more inequality, not less.
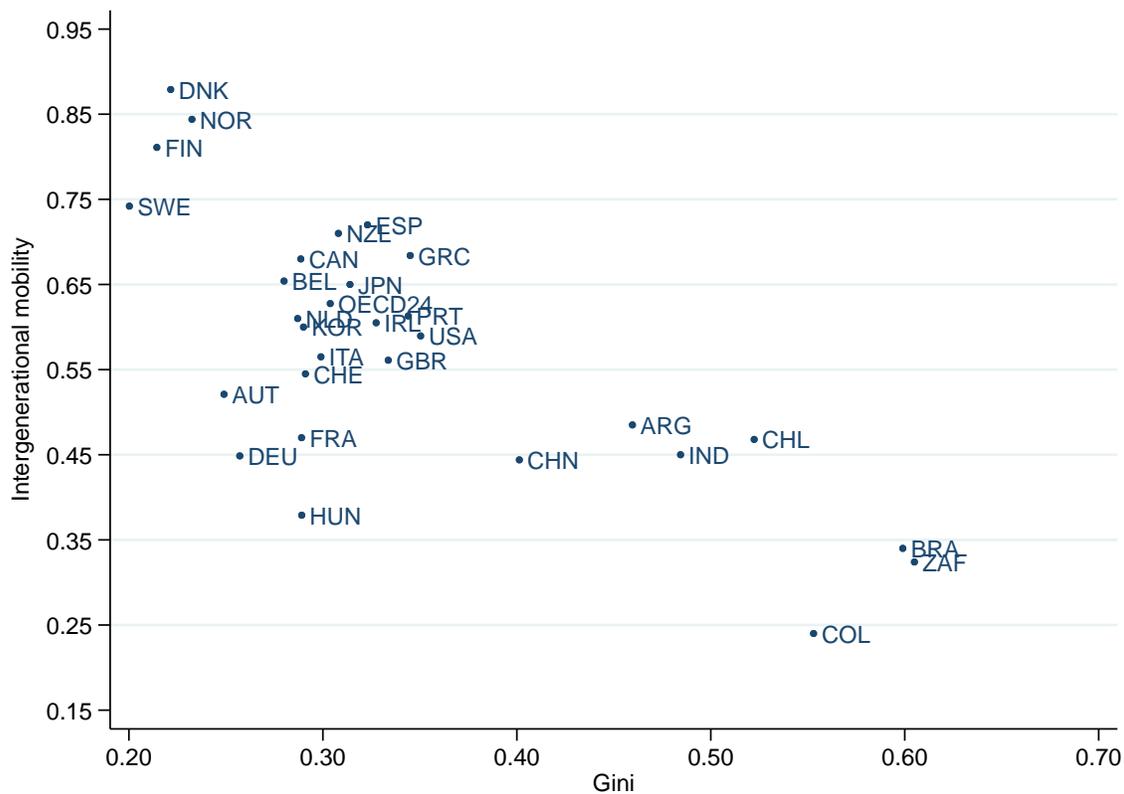
Note: Intergenerational mobility is measured as $(1-\beta)$ and the Gini coefficient is measured from the mid 1980s until the early 1990s. Data Source OECD (2018)

**Figure 1.** Great Gatsby Curve

In any case, using cross country data to estimate the association between variables is problematic since it is difficult to account for relevant heterogeneity. When intergenerational mobility is one of the variables of interest one faces the additional problem that the estimates are often based on limited data that suffer from measurement errors. The gravity of these problems may vary across countries. Another challenge is that some countries may be in a period of transition while others could be in a steady state. Karlson and Landersø (2021) illustrate this point in their comparison of intergenerational mobility in educational attainment in Denmark and the United States. They argue that heterogeneity in intergenerational educational mobility may simply reflect that these economies are, at the time of comparison, at different stages in a development process towards a highly educated modern society. Figure 2 depicts the regression coefficients of children's education on parents' education over a span of 70 years for Denmark and the U.S.. For the U.S. the mobility estimates are relatively stable over the entire 20th century. For Denmark, in contrast, there has been a development from

15

Educational mobility

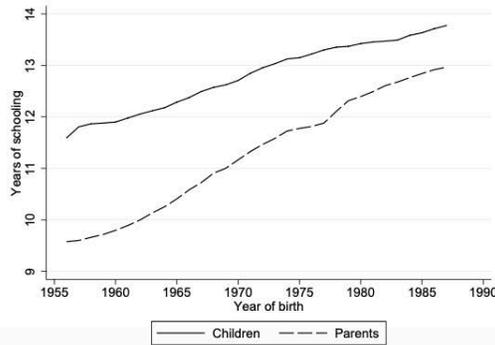a) Children's years of schooling regressed on parents' years of schooling
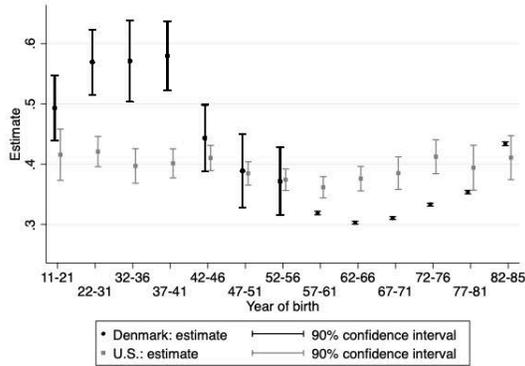
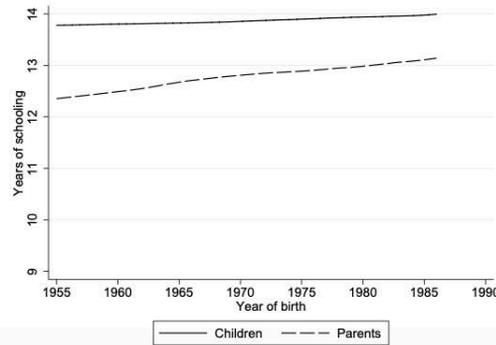b) Upward mobility, children's years of schooling higher than parents' years of schooling
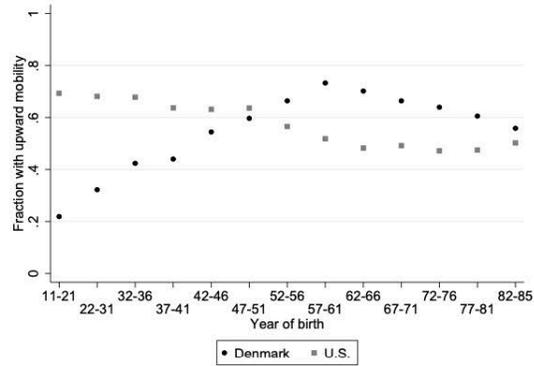
Average years of schooling

c) Denmark

d) U.S.

Note: This is Figure 1 in Karlson and Landersø (2021)

Figure 2. Education IGE in Denmark and the U.S. in the 20th century

very low intergenerational mobility before the second world war, to very high intergenerational mobility in the 50, 60 and early 70s, and a decline in mobility again in the late 70s and the 80s. Karlson and Landersø (2021) argue that the Danish increase in mobility for cohorts born in the early 40s to the mid-60s stems from a reform that expanded lower secondary schooling, which dramatically reduced the education gap between rural and urban Denmark. This period of transition came earlier in the U.S., which means that the period when educational mobility is much higher in Denmark than in the U.S. is a period of educational transition in Denmark, while the U.S. is closer to a steady state.

16

Administrative data sources

Data limitations have constrained our empirical knowledge of intergenerational earnings or income persistence. For a long time the Nordic countries were the shining exception, where researchers have been able to access administrative data for a few decades. More recently, other countries have followed suit and given researchers access to administrative data.[6] Key advantages of such data sets are the accuracy of the income information provided, the large sample size, and the lack of attrition for reasons other than migration or death.

Pekkarinen et al. (2009) is an early study that took advantage of the large sample sizes in administrative data to estimate regional variations in IGE within a country. They studied how a comprehensive Finnish school reform that abolished early school tracking influenced the degree of intergenerational mobility in Finland. The reform had a staggered regional implementation that Pekkarinen et al. (2009) exploited through a difference-in-differences approach. Later, Chetty et al. (2014a) use administrative data to estimate spatial variation in intergenerational mobility across areas within the US. They also explore how mobility correlates with various area characteristics. One of their findings is that areas with more inequality, as measured by Gini coefficients, have less mobility. This pattern is also found in Canada in Corak (2020) as well as in several other developed countries (Stuhler, 2018). As emphasized by Chetty et al. (2014a), the observed geographic variation in intergenerational mobility does not necessarily reflect causal effects of neighborhoods. It could simply be due to omitted variables, such as the types of people living in an area. We discuss this in greater detail in Section 6.

Another advantage with large administrative data sets is that researchers can obtain a more detailed picture of how intergenerational mobility varies across the parental income distribution. Policymakers and researchers are especially concerned about low mobility in both tails of the income distribution. A sticky floor may suggest that children from the lower end of the income distribution either lack talent or do not get the opportunities to realize their potential. A sticky ceiling might be a sign that upper class children inherit remarkable talent or are offered a lot of opportunities. While it can be difficult to differentiate between these explanations, large administrative data make it possible to at least depict and analyze non-linearities in intergeneretional mobility (see e.g. Bratsberg et al. (2007); Bratberg et al.

---

[6]The review article of Røed and Raaum (2003) offers an early and insightful discussion of the many advantages (and some disadvantages) to administrative data. This paper also contains early examples of studies from various countries that have used administrative data to answer important questions. These examples serve as a reminder that administrative data, such as tax records, have a long tradition in empirical economics in general and in research on intergenerational mobility in particular.

(2017); Helsø (2021); Björklund and Jäntti (2020)).

While administrative data have several advantages, it is important to recognize that they are no silver bullet. The fundamental problem of selection bias remains, and even statistical inference may still be an issue, as illustrated in our discussion in Section 6 of neighborhood effects and intergenerational mobility. In addition, administrative data often come with their own measurement problems. The data are collected for government purposes (e.g. taxation), which is reflected in the type (and often limited number) of variables that are recorded and how these variables are measured.

4   The importance of family environment and genetics

While it is important to document intergenerational persistence of socio-economic outcomes, the ultimate goal is to reach beneath the surface and understand *why* family background matters for schooling, income and wealth. In this section we review studies that estimate how much of the variation in such outcomes can be explained by genetics versus family environment. First we discuss the use of sibling correlations in outcomes to obtain an omnibus measure of the role that family background plays. Sibling correlations contain the influence of both genetics and family environment. Next, we describe empirical strategies used to separate the influence of genetics and family environment, and discuss some key findings and their economic interpretation and policy relevance.

4.1   Sibling correlations

Sibling correlations are frequently used to construct an omnibus measure of how family background affects children's income or education. These correlations reflect not only the impact of shared genes but any shared family environment. The basic idea is that the correlation in economic outcomes between siblings will be low if family background plays a minor role for how well individuals do in the economy.

To be precise, it is useful to express the earnings (or any other life outcome) for individual i that was raised in family j as $Y_{ij} = a_j + b_{ij}$, where a captures the family component shared by all siblings and b is the sibling specific component. Since these components are by construction independent, the share of the variance in earnings that is explained by the family component is given by $\rho^Y_{sib} = \frac{\sigma^a}{\sigma^a + \sigma^b}$.[7]

---

[7]In an early study of sibling correlations Solon et al. (1991) show that transitory shocks to earnings will (just as for the estimation of IGE) attenuate the degree of sibling correlation in permanent income.

It is useful to observe that sibling correlations necessarily explain more of the variation in offspring earnings (or in any other life outcome) than do parents' earnings alone. This is simply because one of the family components in *a* is parents' earnings, and *a* also contains other family and neighborhood factors that are relevant for offspring's earnings, but independent of parents' earnings. Solon (1999) show that the sibling correlations and intergenerational correlations (IGC) can be linked in the following way:

$$\rho_{sib}^{Y} = (\text{IGC})^2 + \text{ all sibling shared factors not correlated with parent } Y. \tag{5}$$

This expression is useful for two reasons. It helps us interpret and compare sibling correlations and intergenerational correlations, and it allows us to construct bounds on quantities of interest, as discussed in Björklund and Jäntti (2020).

We can derive a lower bound on family influence from the correlation in outcomes between siblings. It is a lower bound because a given family may affect children differently. For example, siblings may get different genetic endowments from their parents, the birth order may matter for outcomes, and there could be temporal changes in the family environment that will create differences between siblings. Sibling correlations will erroneously assign all these factors—non-shared genes and family factors that affect siblings differently—to the individual component *b*, not the family component *a*.

We can construct an upper bound on family influence from the correlation in outcomes among monozygotic twins. The argument is that monozygotic twins get the same genetic endowment from their parents and, since they are born at the same time, they also share the entire family environment. However, to the extent that monozygotic twins are treated more equally by the environment than ordinary siblings, and influence each other more than ordinary siblings do, their resemblance in outcomes could overstate how important family background is for the population at large.

In Table 1, we report sibling and twin correlations from a set of empirical studies. The sibling correlations are much higher for monozygotic twins than for ordinary siblings. It is also interesting to observe that the cross-country patterns in Table 1 are broadly similar to those reported for the IGE. Specifically, sibling correlations indicate that family background is more important for educational attainment and earnings (income) in the U.S. than in the Nordic countries. By comparing the IGC of an outcome Y, for example educational attainment or earnings, with sibling correlations in Y one can use the expression in equation (5) to calculate how much of the variation in Y^{child} is explained by the variation in Y^{parent}. Björklund and Jäntti (2020) make this comparison and conclude that the IGC in education and

Table 1. Sibling and twin correlations from a set of empirical studies

| Country | Study | Outcome | Ordinary Siblings | | | Twins | |
			Brothers	Sisters	Mixed	MZ-male	MZ-female
Denmark	(Schnitzlein, 2014)	long run earnings	.20	.19			
Germany	(Schnitzlein, 2014)	long run earnings	.43	.39			
USA	(Schnitzlein, 2014)	long run earnings	.45	.29			
USA	(Mazumder, 2008)	long run earnings	.49				
Norway	(Pekkarinen et al., 2017)	long run earnings	.32				
Sweden	(Björklund et al., 2009b)	long run earnings	.35				
Norway	(Björklund et al., 2009a)	years of schooling			.41		
Denmark	(Bredtmann and Smith, 2018)	years of schooling	.31	.39	.33		
Sweden	(Björklund and Jäntti, 2012)	years of schooling	.43	.40			
USA	(Mazumder, 2008)	years of schooling			.60		
Finland	(Hyytinen et al., 2019)	long run earnings				.54	.41
Sweden	(Benjamin et al., 2012)	long run income				.62	.48
Sweden	(Björklund and Jäntti, 2012)	long run earnings				.73	
Sweden	(Björklund and Jäntti, 2012)	years of schooling				.75	.73
Australia	(Miller et al., 2001)	years of schooling				.65	.70
USA	(Behrman and Taubman, 1989)	years of schooling				.75	

Note: This table is based on Table 1 and Table 2 in (Björklund and Jäntti, 2020)

income explains relatively little, roughly ten percent, of the sibling correlations in these outcomes. This suggests that factors other than parental income and education are likely to be important for the observed differences across children in income and education.

The importance of genes

We now discuss two strands of the literature that try to separate genes from family environment. One uses standard methods in behavioral genetics to isolate the degree of heritability. The other uses quasi-random assignment of adopted children to estimate the causal impact of growing up in one family environment versus another.

Heritability and the ACE model

The canonical model for inferring the importance of genes relative to family environment is the so-called ACE model. The basic version of the ACE model assumes that the phenotype, educational attainment or earnings for example, for individual i in family j (Y_{ij}) can

be represented by an additive function of genes (A_i), shared family environment (C_j) and idiosyncratic influences (E_i):

$$Y_{ij} = hA_i + bC_j + dE_i. \tag{6}$$

If the genetic component is independent of family environment, the degree of heritability, which is defined as the fraction of the overall variance in the phenotype that can be attributed to the genetic component, is given by $h^2 = \frac{VAR(A_i)}{VAR(Y_{ij})}$. By comparison, the contribution of family environment is given by $c^2 = \frac{VAR(C_i)}{VAR(Y_{ij})}$.

Heritability can be estimated by comparing the correlation in outcomes for sibling pairs that differ in how genetically related they are. If we normalize Y with a standard deviation of 1, the correlation in outcomes for a sibling pair of type k is given by $\rho_{YY'}^k = \rho_{AA'}^k h^2 + \rho_{CC'}^k c^2$. The difference in correlation between monozygotic (m) and dizygotic (d) twins is then given by $\rho_{YY'}^m - \rho_{YY'}^d = \left(\rho_{AA'}^m - \rho_{AA'}^d\right) h^2 + \left(\rho_{CC'}^m - \rho_{CC'}^d\right) c^2$. If we assume that monozygotic twins share 100% of their genes while dizygotic twins share 50% of their genes, but both types of twins have the same degree of shared environment, we get $h^2 = 2\left(\rho_{YY'}^m - \rho_{YY'}^d\right)$. Comparing ordinary siblings and adopted siblings gives similar expressions (see e.g. the discussion in Sacerdote (2007)).

Findings from the ACE model

According to a recent meta-study by Polderman et al. (2015), over the last fifty years 2,748 publications have used nearly 15 million pairs of twins to estimate the heritability of 17,804 human traits. The average heritability for all traits tested is around 0.5 and physical traits such as height have higher heritability (around 0.8) than more complex behavioral outcomes. Taubman (1976) is an early study that uses twins to infer the genetic component of earnings. Since then, this framework has been used extensively by social scientists to estimate the heritability of many socio-economic outcomes, most often educational attainment, earnings, and income, but more recently also wealth and other aspects of the economic phenotype, for example risk preferences and entrepreneurship (Nicolaou et al., 2008).

Hyytinen et al. (2019) use long panels of earnings for monozygotic and dizygotic twins to estimate a heritability of lifetime earnings of 53% for men and 39% for women (see Table 1). Their findings are broadly in line with other studies of the heritability of earnings and income, such as Sacerdote (2011) and Björklund and Jäntti (2020). In a meta study, Branigan et al. (2013) find that, on average, around 40% of the variance in educational attainment can be attributed to variation in genetic components. The heritability of educational attainment will naturally vary across countries, depending on environmental factors such as access to

and quality of educational institutions. Engzell and Tropf (2019) combine cross country data on intergenerational mobility in education and data from twin studies of the heritability of education. They find a positive association between heritability and intergenerational schooling mobility. A possible explanation is that in a society with equal opportunities and universal access to higher education, mobility will be relatively high. Since everyone has equal opportunities to choose higher education, ability will explain much of the variation in education and, as a result, heritability will be high (Trzaskowski et al., 2014).

Limitations and critiques of the ACE model

The basic ACE model rests on several restrictive assumptions. One assumption is that ordinary siblings share, on average, 50% of their genes. However, that number is probably too low. Individuals tend to mate and have children with persons who resemble themselves (Kalmijn, 1998). With assortative mating, dizygotic twins or siblings more generally will (in expectation) share more than 50% of their genes and, as a result, the ACE model underestimates the heritability of traits.

Another restrictive assumption of the ACE model is that it assumes independence between genes and the environment. It also assumes no interaction effects between genes and the environment. For the environmental factors and genes that produce the economic phenotype, it seems likely that good genes are correlated with a good environment, partly because the genes "choose" and shape their environment (Plomin et al., 1977). There is also an increasing body of evidence, discussed below, suggesting that the impact of genes, the way they express themselves in a phenotype, depends on the environment.[8]

The ACE model further assumes that all sibling types, irrespective of their genetic relatedness, share family environment to the same degree. This is a strong assumption. It is likely that monozygotic twins are treated more equally by parents and peers than dizygotic twins, and, as a consequence, monozygotic twins may share more environmental factors than dizygotic twins. Some of the additional correlation between monozygotic-twins could therefore come from their extra shared environment.

When assessing the restrictions of the basic ACE model, it is useful to observe that additional data may allow one to relax some (but not all) of its strong assumptions. This is possible if the analyst has access to data on a number of different types of sibling pairs such

---

[8]Epigenetics studies how the environment impacts how genes express themselves in phenotypes. The fact that the genome contains environmentally influenced epigenes that regulate the behavior of genes makes the dichotomy between nurture versus nature misleading. It is more appropriate to talk about nature via nurture; for an interesting layman's introduction to these ideas, see Ridley 2003.

as monozygotic twins, dizygotic twins, regular siblings, half siblings and adoptees. Data like this has been used to allow for gene environment correlations and differences in the degree of shared environment across siblings (see e.g. Bjorklund et al. (2005) and Fagereng et al. (2021)). However, even in these cases the results from the ACE model need to be interpreted with caution.

Genes as covariates

The genetic component *A* in the ACE model is a latent variable, a black box. Its role in explaining variation in some outcome can only be inferred by making strong independence and functional form assumptions about how genes and the environment produce a phenotype. Since the start of this century it has been possible to crack open the black box and map the whole human genome. This possibility has transformed the study of how genetic variation affects human traits and outcomes. It is now common to use data driven estimation procedures to search for single letters in the genetic code—single-nucleotide polymorphisms (SNPs)—that correlate with a phenotype of interest.

Research that uses this approach is known as genome-wide association studies (GWAS). To summarize how genetically predisposed a person is to develop a trait, or obtain an outcome such as years of schooling, it is common to create a summary measure, a risk score, by weighting the genome of a person with the GWAS estimated SNPs coefficients (either all SNPs are used, or only those that are significantly related to the trait). This score is denoted the polygenic score (PGS).

In a sample of 293,723 individuals, Okbay et al. (2016) find 74 SNPs that have a genome-wide statistical significant association with educational attainment. Lee et al. (2018) conduct genetic association analysis of educational attainment on a larger sample of 1,1 million individuals and find 1,271 genome-wide-significant SNPs. The SNPs that are correlated with educational attainment are disproportionately located on genes that are involved in brain development and neuron-to-neuron communication. They construct a PGS for educational attainment based on on their data, denoted the EA-score, and apply this score to different data sets. In these cross validations, the EA-score explains 11–13 % of the variance in educational attainment. The explanatory power of the EA-score is considerably higher than earlier PGS scores based on smaller GWAS samples, but still quite a bit below the the fraction of variation in educational attainment that is attributed to genetic components in the ACE model. This gap between ACE estimates of heritability and the R^2 based on PGS is commonly referred to as the missing heritability problem. There are several reasons why models that use poly-

genic scores as covariates might explain less of the variation in educational attainment (or any other outcome) than ACE models. One explanation could be that the assumptions in the ACE model do not hold and it gives a biased estimate of heritability. However, according to Cesarini and Visscher (2017) there is now a broad agreement in the literature that heritability may be hiding rather than missing in PGS models. The point is that many relevant SNPs have such a low impact on an outcome that they are not detected in the sample sizes usually applied in these studies.

Besides the missing heritability problem, there are also important questions and concerns about how to interpret the findings of studies that use genes as covariates. One issue is that the SNPs that are correlated with educational attainment do not necessarily capture the effect of changing the genetic information at these locations. Suppose individuals in generation t with a certain trait totally irrelevant for educational aptitude (long nose), were for some reason favored in the educational system. Suppose that in generation $t + 1$ access to education is a leveled playing field. If we conducted a GWAS on educational attainment in the $t + 1$ population, we would nevertheless find that SNPs coding for a long nose would be associated with high education. This is because long-nosed parents give their kids long noses through genetic transmission, but affect the education of their children through environmental transmission, for example by having the means to invest in children's human capital. Thus, even if we blocked the genetic channel, long-nosed parents would have more educated children.

Another mechanism that makes it difficult to separate nature from nurture is known as genetic nurture (Kong et al., 2018). It is possible that the SNPs that code for educational aptitude/attainment will have a positive effect on the education of offspring, even if the relevant alleles are not transmitted to the offspring. The reason is that parents who are genetically predisposed to education may change the environment of their offspring to stimulate schooling.

Finally, it is worth emphasizing that the quantitative importance of PGS scores may be sensitive to environmental factors. For example, Papageorge and Thom (2020) apply EA-scores to a new data set, the Health and Retirement Study. They find that the EA-score predicts a higher degree of graduation and the score explains around 7.5% of the variation in years of education, which is a bit lower than the findings in Lee et al. (2018). Interestingly, they find that the relationship is not very strong for families with a low socioeconomic score. This result agrees with the so called Scarr-Rowe hypothesis that genetic variation plays less of a role in socially and economically deprived families. They also find that variation in EA-score explains some of the variation in earnings even after accounting for the level of education. Barth et al. (2020) show that the individual EA-score is strongly related to wealth at retirement. They show that the relationship between EA-score and wealth remains even if

income and education are included as explanatory variables.

The policy relevance of genoeconomics

In the social sciences, the heritability of human traits and achievements has been a highly debated and controversial topic. Part of the controversy comes from the fact that the concept is often misunderstood. A heritability estimate measures the fraction of the population variation in individual outcomes, such as education or income, that is explained by genetic variation in the population. The degree of heritability of an outcome cannot tell us how important genes are for shaping the outcome, nor how easy it is to change the outcome. A natural and relevant question is whether, and in what situations, heritability estimates can be useful for economic policy.

Manski (2011) discusses several objections to the policy relevance of heritability estimates from the ACE model. One objection is that the nature of policy interventions is to change the environment, while heritability is calculated from data collected in the environment that prevailed before the intervention. For example, the heritability of educational attainment will depend on the amount of heterogeneity in the quality of primary schools in the population. If we changed the environment and made schools more unequal, this would reduce the the heritability of educational attainment. This is a valid point, but it is not specific to empirical research on heritability. It is a general concern in empirical analysis that a parameter estimated on data from one population or in one environment may not generalize to other populations or to other environments.[9] To gauge the degree of invariance in the heritability of educational attainment, it is possible to empirically examine or model how the parameter would differ in an alternative or counterfactual environment.

A second objection is that a high degree of heritability does not, in itself, tell us anything about the effectiveness of policy interventions, such as how costly it is to alter outcomes. As Goldberger (1979) pointed out, the heritability of bad eyesight is very high but can readily be fixed by good opticians. Again, this critique is not specific to heritability estimates. The observation that the root cause of a problem may not be relevant for how to best solve it is a general insight. For example, the best way to avoid getting wet if it rains could be to stay indoors or bring an umbrella, not to change the rainy weather. And the most effective way to reduce labor market inequalities could be to change the tax-transfer system, even if one thinks globalization and technological changes are the root causes of increased labor market

---

[9]See, for example, Heckman (2005) for a broad and insightful discussion of structural models, treatment effects, and invariance assumptions in econometrics.

inequalities.

Finally, it is worth stressing that even if the bulk of the variation in an outcome is explained by nature, it does not imply that it is natural in the normative sense that it ought to be like this. A high degree of heritability in an outcome has no implication for whether we ought to reduce the outcomes differences across individuals. To what extent a society should have more or less equal outcomes in a particular dimension is a separate, normative question—and one that economics as a science has little, if anything, to say about.

Manski (2011) also discusses the policy relevance of the GWAS approach, concluding that it is potentially more useful. With GWAS, genetic information can be used as covariates in policy analysis to understand and predict heterogeneity in treatment effects. Schmitz and Conley (2017) provides an example of this: they interact an instrument for educational attainment (the Vietnam draft lottery) with EA-score and find that those with a low EA-score lost more education due to being drafted. While heterogeity in treatment effects may be of interest to policymakers, the arguments against the policy relevance of heritability estimates apply here as well. For example, the predictive power of a PGS-score is estimated for a particular pool of genes and for a particular educational system, and may not be valid in a different population or in a reformed education system.

Irrespective of this caveat, genoeconomics has the potential to increase our understanding of why individuals respond differently to policy changes. Can it do more? In principle, PGS for educational attainment, such as EA-score or individual markers for SNPs that are highly associated with educational attainment, could be used to individualize school policy based on an individual's genetic endowments. In their assessment of the practical relevance of GWAS, Cesarini and Visscher (2017) characterize this idea as highly speculative. This use of genetic information would also raise ethical concerns. At a more practical level they foresee that polygenic scores can be useful as controls in randomized intervention studies to reduce residual variance and increase precision. However, it remains to be seen if such scores are going to be empirically important in explaining or predicting individual responses to interventions or policy.

4.2 *The impact of family enviornment*

As alternative approach to the ACE model for gauging the impact of the family environment is to vary the environment while holding the genetic relatedness to parents constant. The ideal experiment would be to randomly assign newborn children to parents who have different education, income and wealth and who live in different neighborhoods. To see what is (not)

possible to identify with such an experiment (this is discussed in greater detail in Holmlund et al. (2011); Fagereng et al. (2021)), consider the extended intergenerational transmission equation

$$Y_i = \alpha + \beta Y_{j(i)} + X'_{j(i)} \eta + \gamma \kappa_{j(i)} + X'_i \lambda + \delta \chi_i + \varepsilon_i, \tag{7}$$

where Y_i is the outcome of interest of the child i, say earnings; the characteristics of her family j consist of parental log earnings $Y_{j(i)}$ and a vector of observable (pre-determined) family characteristics other than earnings $X'_{j(i)}$ and an unobservable component $\kappa_{j(i)}$. Similarly, the child has (pre-determined) observable characteristics X'_i (e.g. gender and birth cohort) and unobservable characteristics χ_i such as genes. The scalar error term ε_i is constructed to be orthogonal to all other variables in the equation. The unobservable variables that may correlate with the explanatory variable of interest, parental income, are κ_{ji} and χ_i.

With random assignment of children to families, the potential outcomes, defined by the genes of a child, are uncorrelated with the family environment the child grew up in, that is, χ_i is independent of the family components $Y_{j(i)}, X'_{j(i)}, \kappa_{j(i)}$. Random assignment of children, therefore, makes it possible to estimate the causal effect of being raised in a high earnings family versus a low earnings family. One cannot, however, without making further assumptions, use random assignment of children to estimate the ceteris paribus effect that an exogenous increase in the family's earnings would have on the child's outcome. There is likely to be correlation between $Y_{j(i)}$ and $\kappa_{j(i)}$: higher earning families may have other unobservable qualities in their family environment that also affect the child's outcome. As discussed in Section 5, drawing causal inference about the ceteris paribus effect of an exogenous increase in parental earnings requires random variation in the earnings of a given family, not random assignment of children to a high earnings family versus a low earnings family.

Of course, most children grow up with their biological parents, but some do not, and comparing outcomes of adopted children raised in different families is a frequently used empirical strategy to gauge how different family environments affect children's outcomes. However, to interpret differences in outcomes among adopted children as the causal effect of a family environment, the children must be randomly assigned to families. Kinship adoption is relatively common in many countries, and is clearly at odds with the notion of random assignment of adoptees to families. Even for non-relative adoptions, there is likely be a genetic association between child and parents in settings with selective placement of adopted children, either based on requests from adopting parents or from matching criteria by adoption agencies.

These concerns are the reasons why Sacerdote (2007) and Fagereng et al. (2021) use

data from infant adoptees from Korea to the US and Norway, respectively. Both substantiate quasi random assignment of these adoptees to pre-approved adoptive families, by providing detailed descriptions of the placement rules, and by checking that observable features of the adopting family do not predict pre-adoption characteristics of the adoptees. Sacerdote (2007) study several outcomes, among them income and education. He finds no effect of being assigned to a higher earning family, but adoptees who were assigned to small families in which the mother has higher education tend to have higher educational attainment themselves. He also find strong family environmental effects on smoking and drinking habits of the children as adults.

Fagereng et al. (2021) estimate how wealth and financial risk taking are related between parents and their randomly assigned adopted children. They find that children who were assigned to wealthier families are significantly richer in adulthood. On average adoptees accrue an extra US$2,250 of wealth if assigned to an adoptive family with US$10,000 additional wealth. The magnitude of this estimate suggests that adoptees raised by parents with a wealth level that is 10% above the mean in the parent generation can expect to obtain a wealth level that is almost 3.7% above average for their own generation. They also find that adoptees' stock market participation and portfolio risk are increasing in the financial risk position of their adoptive parents. To interpret the importance of family environment in wealth transmission they compare the intergenerational transmission association in wealth for adoptees with non-adopted children. They find that the predictive influence of parental wealth on children's wealth is twice as large for biological children as for adoptees.

Several other studies use the outcomes of adoptees to address the question of nature versus nurture for children's outcomes; see Holmlund et al. (2011) for an overview of the literature. None of these studies, however, can substantiate that adoptees are randomly assigned to parents. In fact it is generally acknowledged by the authors that adoptions are selective, and with selective placement it is difficult to separate the influence of genes from the influence of the family environment. One indirect solution is to find a proxy to control for the genetic disposition of the adopted child. This is the approach taken by Björklund et al. (2006). They use data from Swedish adoptees and are able to observe income and education (at least partially) for both the rearing parents and the biological parents of the adoptees.

In their study, Björklund et al. (2006) regress education of the adopted child on the education of all four types of parents. They find the following transmission coefficients for years of schooling: 0.13 for biological mother, 0.11 for biological father, 0.11 for rearing father, 0.07 for rearing mother. Interestingly, the sum of the coefficients for the biological and the rearing mother of the adoptees resemble the coefficient for the standard educational

IGE on the Swedish data. The same is true for the biological and the rearing fathers. This points to the possibility of a simple additive structure of the influence of pre-birth (nature) and post-birth factors (family environment) on the outcomes of children in a variety of family types, a structure that is explored further in Bjorklund et al. (2005). Black et al. (2020) use the same model to estimate family environment effects on wealth transmission using Swedish adoptions. They also find substantial family environmental effects on wealth transmission. The critical assumption in this approach is that selection bias in non-random adoptions is adjusted for by controlling for the observed outcomes of biological parents. A natural concern is that the observed outcomes (e.g. wealth) of parents giving their child up for adoption may, in part, reflect adverse shocks and, thus, be poor control variables for biological parents potential outcomes (e.g. wealth).

A research design with quasi-random assignment of adoptees to families, as in Sacerdote (2007) and in Fagereng et al. (2021), has strong internal validity: random assignment identifies the treatment effects of being raised in different families within the sample of adoptees. Nevertheless, a natural question is to what extent the effects based on random adoptees generalize to the population of children at large. There are two reasons that may limit the external validity of these studies. The parents who adopt may be and behave different from non-adopting parents, and therefore also influence their children differently, or the adopted children may be different from non-adopted children. Fagereng et al. (2021) examine these threats to external validity in great detail. One check they do is to estimate the intergenerational transmission of wealth within the subsample of adoptive parents with both biological and adopted children. They find that the difference in wealth transmission between biological children and adoptees within this sample is roughly the same as the difference they find when comparing biological and adoptees that have different parents. This indicates that the parents who adopt in their data are similar to other parents when it comes to intergenerational wealth transmission. By comparison, external validity could be more limited in studies that use data from non-random domestic adoptions. When Fagereng et al. (2021) make a comparison of adopting and non-adopting parents for domestic adoptions in Norway, they find substantial differences. Hence, analysis based on non-random domestic adoptions may lack both internal and external validity.

5 Effects of family environment

The research to date suggests that less than half of the population variance in education, earnings and wealth is explained by genes. This means that variation in nurture—variation in

environmental factors—plays a key role in explaining individual life outcomes. In this section we review studies that attempt to quantify intergenerational causal effects of specific family environmental factors. We first review studies that estimate the extent to which exogenous changes in parental education and income are transmitted to children. Next, we consider how a broader set of family changes affect children. However, before we turn to the empirical literature, it is useful to briefly consider how changes in parental education or income affect children's outcomes in the theoretical framework outlined in section 3.

Theoretical predictions

In the Becker and Tomes (1986) model, higher parental education has no direct effect on the child, but it may have an indirect effect via earnings.[10] Higher education implies higher income and higher income may induce parents to invest more in their child's education. This transmission channel is only relevant if parents are credit constrained. If parents have access to sufficient internal or external resources, the amount of human capital investment in the child is determined by the returns to the investment and the interest rate, and, therefore, independent of a ceteris paribus change in parental income.

The extent to which parents' investment in human capital is influenced by credit constraints is likely to vary across contexts, depending on how education is financed, on access to student loans, and so on. There is an emerging body of evidence suggesting that credit constraints are indeed important for educational attainment (see e.g. Lochner and Monge-Naranjo (2016); Hai and Heckman (2017)). Credit constraints seem to be particularly salient for parental investment in schooling in developing countries (Attanasio and Kaufmann, 2009; Solis, 2017).

In the extended Beckerian framework with multiple periods of learning and human capital investments and multiple skills, parents play a more active role in shaping the aptitudes and attitudes of their children. There are large differences in parenting styles and child involvement, potentially influencing the extent to which parents motivate and help their kids to reach their potential. An exogenous change in parents' education or earnings may affect their children's outcomes, for example through how much, and how, they interact with their child. (Collins et al., 2000). With complementarities in the returns to human capital investments at different stages of the life cycle, even small improvements in early childhood may facilitate later learning and create real changes in how well these individuals do in adulthood. (Cunha

---

[10]In Becker et al. (2018) there is a direct intergenerational human capital link since they assume that the returns to human capital investments in children increase in the human capital of the parents.

and Heckman, 2007; Heckman, 2007).

Does parents' schooling have an effect on the schooling of their child?

(Holmlund et al., 2011; Björklund and Salvanes, 2011) review the literature examining intergenerational transmission of education. Most of the empirical work uses one of three strategies to identify causal effect of parental education on children's education: (i) adoptions (ii) twins and (iii) instruments for parental education (e.g. school reforms).

We have already addressed identification issues with the first strategy using data on adoptions. With random adoptions, it is possible to estimate environmental treatment effects, for example the effect on adopted children's education of being assigned to parents with high education rather than low education. However, since a parent's education typically correlates with other unobserved characteristics (e.g. parenting style and child rearing skills), this strategy does not, without strong assumptions, identify the intergenerational effects of exogenous changes in parental education.

Another strategy is to consider twins who become mothers (or fathers). If monozygotic twin mothers have different levels of education, their children obtain the same genetic endowment from their mothers, but have been raised by mothers with different schooling. Behrman and Rosenzweig (2002) use this empirical strategy to estimate the causal effect of parental education. They find a small positive effect of fathers' education on offspring, whereas mothers' education has no effect on offspring education. The identifying assumption here is that the difference in twin education is independent of other factors that may affect children's education. This assumption is questionable for several reasons, as discussed in detail in Bound and Solon (1999) and Holmlund et al. (2011).

The third strategy is to find instruments that exogenously shift parental education (without directly affecting children's education). Black et al. (2005) use a staggered implementation of a school reform in Norway that increased basic education from 7 to 9 years to estimate the effect on children's education. The reform did have a positive effect on parents' education, but they find no significant effect of mother's years of education on the schooling of their child. Holmlund et al. (2011) use a similar Swedish schooling reform and they find a positive but modest causal effect of parental education on children's education. A year of extra schooling for mothers (fathers) increased the schooling of the child by 0.09 (0.11) years. A useful feature of Holmlund et al. (2011) is that they compare the results from applying all three empirical strategies mentioned above on the same data set. They find that all methods provide lower estimates than the OLS estimates of intergenerational educational

transmission, pointing to selection bias or heterogeneity. Furthermore, they show that the three methods produce quite different point estimates. This could be because the underlying assumptions fail in one or more of the methods, or because each method identifies the effects for a different subgroup of the population (Holmlund et al., 2011).

Our reading of the literature is that studies that use plausibly exogenous variation in parental education produce much smaller estimates of the IGE in education compared to the OLS estimates.[11] This is also in line with the findings in De Haan (2011), who uses data from the Wisconsin Longitudinal Study to construct a nonparametric bounds analysis. This approach relies on the assumption that mother's education and "parenting quality" influence the schooling outcome for the child in the same direction. She finds that the lower bound is positive and the upper bound is substantially lower than the OLS estimate.

The causal effects of parental income

Most studies of the causal effects of parental income use one of the three types of empirical strategies discussed above. The only exception is that the focus now is on parental income and, thus, the third strategy is based on reforms or natural experiments that generate plausibly exogenous variation in parental income.

A few studies have used data on adoptees or twins in an attempt to draw causal inferences about the impact of parental income. Amin et al. (2011) examine the son-father association in income in a sample of twin fathers. Björklund et al. (2006) use data on adopted children in Sweden. Both studies estimate positive but modest effects of parental income. The estimated effects are about half the size of the usual OLS estimates of the IGE in income.

Cooper and Stewart (2017) review studies from developed countries that examine the impact of family income and resources on children's health, education and social and behavioral development. In total 67 studies have an empirical design that addresses selection on unobservables through an explicit source of (quasi) experimental variation in parental income. A majority of the studies use data from the U.S., and many use variation in EITC as a source of income variation (see e.g. Dahl and Lochner (2012)). Others use lottery prizes or regional income shocks to estimate how changes in household income affect children's outcomes. A vast majority (45 out of 61) of the studies reviewed in Cooper and Stewart (2017) find a significant positive effect of income across a range of children's outcomes, especially for children from low income or disadvantaged families. Due to data limitations, only a few of

---

[11]In their comprehensive review of intergenerational mobility in educational attainment Björklund and Salvanes (2011) conclude that at most half, probably less, of the educational IGE is causal.

the studies reviewed consider the effect of parental income on children's income as adults. An exception is Aizer et al. (2016). They estimate the effect of the Mothers' Pension program in the US, finding that the children of accepted applicants obtain a third more education and income compared to children of mothers who were rejected.

Many of the studies reviewed in Cooper and Stewart (2017) use reforms that change the income of families at the lower end of the income distribution, and, therefore, they are silent about the impact of exogenous changes in income among middle or high income families. A study that examines the effect of a positive income shock also for middle and higher income classes is Løken et al. (2012). They use temporal and regional variation in the oil boom in Norway as a source of exogenous income variation. Their nonlinear instrumental variables estimates show an increasing, concave relationship between family income and children's outcomes. In other words, income matters a lot more for children of the poor than for children of the rich.

Two studies stand out in their findings. The first is Cesarini et al. (2016), who find little effect of an exogenous increase in wealth on children outcome. They use administrative data on Swedish lottery players to estimate the causal effect of lottery winnings on children's outcomes. Except for a modest reduction in obesity risk, they find fairly precisely estimated zero effects on many other child outcomes. Another lottery study, Bleakley and Ferrie (2016), compares the outcomes of the children and grandchildren of the winners and losers of the Georgia's Cherokee Land Lottery of 1832. Participation was almost universal and the prizes were large. They find no effects on the life outcomes of the descendants of the winners. It is hard to reconcile these two lottery studies with the positive effects found in other studies. A possible explanation is that the windfall gains from lotteries have different behavioral and transmission effects than increases in ordinary income.

There are also studies that examine to what extent the timing of parents' income matters for children's outcomes. These analyses speak directly to the extended Beckerian model with multiple periods of parental investments and multiple skills that parents can influence by allocating time and resources to their children. A notable example is Carneiro et al. (2021). In their analysis they split childhood into three intervals (ages 0–5, 6–11, 12–17) and examine the outcomes of children who have the same permanent income but different income profiles over these three intervals. They find that children of parents who have a relatively high income in early childhood and low income in middle childhood do better in terms of education and earnings than children of parents with the opposite income profile. One potential issue, discussed in the paper, is that a high income may reflect longer work hours and that parents with high income may substitute more income for less time spent with their children, which

may have a countervailing effect on child development (see Agostinelli and Sorrenti (2018) and Nicoletti et al. (2020)).

Parental leave, childcare and child development

Another source of variation in the time mothers spend with their children comes from changes in the duration of maternity leave. Several studies consider the impact of reforms that increase the parental leave period for mothers on the educational attainment and earnings of children later in life. The evidence is mixed. Carneiro et al. (2015) use a generous extension of fully paid maternity leave in Norway and a difference-in-differences approach to estimate the effect on children's outcomes. They find that the extra time that mothers spend with their children increases both educational attainment and earnings of the children, especially those with less educated mothers. Cools et al. (2015) find that an extension of paternity leave in Norway had a positive effect on children's school performance. Liu and Skans (2010) find no average effect on children of a maternity leave reform in Sweden, but they do find a small positive effect on children with highly educated mothers. Dustmann and Schönberg (2012) consider a German extension of maternity leave and find no effect, or even a small negative effect on the educational attainment of the children of the mothers that were intentionally treated by this reform.[12]

One possible reason for the differences in the effects of extended parental leave is that the consequences of parents spending more time with the child probably depend on who the alternative care taker is. This point is perhaps even more relevant for the effects of expanding child care. Havnes and Mogstad (2011a,b) exploit variation over time and across regions in the exposure of Norwegian children to formal subsidized child care of relatively high quality. They find positive average impacts on the exposed children's long run outcomes, such as their educational attainment and labor market participation. The reform had a very modest impact on female labor force participation, from which they conclude that the expansion of formal subsidized child care did not replace parental care, but informal day care alternatives. If access to high quality informal care is not available or affordable to low income parents, this may explain why children from low socioeconomic families gained more from the introduction of formal child care.

This argument is consistent with the findings of Havnes and Mogstad (2015). They re-examine the expansion of high quality subsidized child care in Norway, finding that the effects

---

[12]See Currie and Almond (2011) and Almond et al. (2018) for a discussion of how various types of changes in family environment during pregnancy or in the period after birth affect may affect children's short and long-run outcomes.

were positive in the lower and middle parts of the earnings distribution of exposed children as adults, and negative in the uppermost part. They complement this analysis with local linear regressions of the child care effects by family income. They find that most of the gains in earnings associated with the child care expansion relate to children of low income parents, whereas upper-class children actually experience a loss in earnings. In line with the differential effects by family income, they estimate that the universal child care program substantially increased intergenerational income mobility.

Herbst (2017) reaches a broadly similar conclusion in his study of the effect of the Lanham Act of 1940 in the U.S., which provided child care to communities with a demonstrated need for war-time child care. He uses a difference-in-differences approach to estimate the effect of this program on maternal employment and children's outcomes. He finds a positive effect on employment and also long-term positive labor market effects on the children that were treated by the reform. The effect is driven by children from disadvantaged families.[13]

The studies discussed above are some of the few that look at the effects of subsidized child care on the long-run labor market outcomes of children. Several other studies have looked at more immediate effects of subsidized or universal child care. Cornelissen et al. 2018 use data from a staggered expansion of childcare in the Lower Saxony region of Germany to estimate the marginal treatment effects of universal child care. They find large positive effects of participation in child care among children from minority families. Felfe and Lalive (2018) use child care expansion data from another region in Germany and find that universal child care improve socio-emotional skills, especially for boys from disadvantaged families.

For younger children, the effects of child care are mixed. Studies have found that for children below the age of two years, being assigned to child care may have negative effects on child development indicators (Herbst, 2013; Fort et al., 2020). The problem here appears to be that this child care is of relatively low quality. Drange and Havnes (2019) use data from a child care lottery in Norway—a lottery to get access to high quality care alternative—and find positive effects on the cognitive development of children who get access when they are 15 months rather than 19 months, and the effects are again largest for disadvantaged families. Their study provides further evidence that high quality child care can increase social mobility.

These studies illustrate that family policy may play an important part in promoting equality of opportunity across socioeconomic classes, either by allowing parents to invest time in

---

[13]These findings are consistent with the positive long-term effects found in high quality, targeted early interventions such as the Perry Preschool program. See Heckman et al. (2010) and Heckman and Karapakula (2019) for a thorough evaluation and discussion of this program. Cascio (2021) contains a comprehensive discussion for the design and evaluations of early child care education policies in the U.S.

their children during early childhood or by offering subsidies to universal high quality institutionalized child care. To have positive effects on intergenerational mobility, it is important that the institutional or formal child care that replaces parental care, or informal day care, is of sufficiently high quality. A caveat is that disadvantaged families may have a high resistance to use high quality child care alternatives, as found in Cornelissen et al. 2018, either because they cannot afford these alternatives or because they are unaware or misinformed about their potential benefits. Boneva and Rauh (2018) study parents' beliefs about how their investments in their children affect their outcomes later in life. They find that parents perception of the returns to investment is associated with their socioeconomic status. Low income households tend to think the returns to early childhood investment are relatively low. In their data these perceptions are associated with actual child investments. If early childhood investments are as important as many other studies indicate, these socioeconomic differences in beliefs may underpin high intergenerational transmission of economic success.

It is, however, not obvious that the stated perceptions and beliefs cause action or inaction. It could be the other way around, that alternatives that are out of reach are deemed unproductive. The findings reported in Boneva and Rauh (2018) are nevertheless intriguing. See also Abbiati and Barone (2017) for similar findings from Italy and Stuhler (2018) for a more complete discussion of beliefs and information frictions and social mobility. These findings point at two important questions for future research: (i) to what extent is it lack of information and misconceptions rather than lack of resources that prevent poor families from making productive investments in the human capital of their children; (ii) if parents have incorrect beliefs with respect to parenting, for example that early and late investments are substitutes rather than compliments and later investments have higher returns than early investments, how can policy interventions correct such beliefs.

6 Neighborhoods and Intergenerational Mobility

The literature on neighborhood effects covers a range of theoretical, econometric and empirical issues. As such, it defies easy summary. Our discussion will be selective, centered around relatively recent work on how neighborhoods may influence intergenerational mobility. For more extensive reviews, see Durlauf (2004) and Galster (2012).

6.1 How Should Neighborhoods be Defined?

A precursor to any analysis of neighborhood effects is a definition of neighborhoods. This requires an answer to the difficult but important question of how to determine neighborhood boundaries for social measurement. Sociologists have long debated this question. In theory, neighbourhoods can be defined as geographically bounded groupings of households and institutions, connected through structures and processes (Coulton et al., 2007). However, the empirical counterpart to this type of neighborhood definition is rather unclear and, as a result, subject to much controversy.

We argue there are at least two key challenges to empirically defining neighborhoods. The first and most fundamental challenge is that the analyst rarely observes the relevant boundaries of neighborhoods. This challenge should come as no surprise to economists, who struggle to determine the boundaries of markets. In theory, a market is the environment within which a price is determined. Or, in other words, a market is the set of traders whose demand and supply establishes the price of a good (Stigler and Sherwin, 1985). Thus, an appropriate definition of a market depends on the elasticities of demand or supply: two goods may be considered in the same market if the cross-price elasticities are sufficiently high, so that changes in relative prices would lead to large changes in the relative quantities produced or purchased. In this spirit, one may argue that two localities should be considered in the same neighbourhood if each location's characteristics (e.g. physical attributes or inhabitants) have a sufficiently strong impact on individuals in the other location. Whether these localities are geographically adjacent is not necessarily important. Individuals in one location may both be exposed to and affected by the physical attributes or inhabitants of a non-adjacent location.[14]

A natural way to empirically assess the boundaries of neighborhoods would be to estimate the cross-location effects of exogenous changes in the characteristics of a given location. Sharkey (2010) takes this approach. He studies the exposure effects of a local homicide on the cognitive performance of children, reporting that the estimated exposure effects decay with distance. For example, a homicide within the child's census track has a larger exposure estimate than a homicide in adjacent census tracks. While rarely performed, this type of cross-location analysis is akin to the usual market test in merger evaluations, where cross-price elasticities of demand are used as a metric for product substitutability and market

---

[14]The analysis of neighborhoods in Conley and Topa (2002) acknowledges this. The authors study 75 community areas in Chicago with the goal of identifying areas with a common sense of community. To this end, they use unemployment data from these community areas and construct spatial correlation functions to understand how the neighborhoods covary along four different notions of distance: physical distance, travel time distance, racial and ethnicity distance, and occupational distance.

boundaries.

The second challenge to defining neighborhoods is that locations consist of a vector of characteristics, including both physical attributes and inhabitants, and each of these characteristics may affect individuals' outcomes in complicated ways. Durlauf (2004) and Galster (2012) provide extensive reviews of the different mechanisms through which various characteristics of neighbourhoods could matter for individuals' outcomes. Galster (2012) classifies neighbourhood characteristics into four groups depending on the type of mechanism though which they operate: social-interactive, environmental, geographical, and institutional mechanisms. For example, environmental mechanisms include measures of local pollution such as air quality, whereas local labor market conditions and access to jobs are examples of geographical mechanisms. While an individual may commute to other locations to work, the air quality to which she is exposed may largely be determined by where she lives. This example illustrates that the appropriate definition of neighbourhoods is not necessarily the same across the four neighbourhood mechanisms.

The curse of dimensionality that naturally arises when defining and analyzing neighborhoods with a wide range of characteristics does not necessarily arise in economic analyses of markets. The different characteristics of a good can often be reduced to a single dimension, a price, and changes in relative prices should determine the behavior of the traders in a market. Of course, prices may also be informative when defining neighborhoods. For example, certain neighborhood characteristics, such as local public services and taxes, should be reflected in property values, possibly creating discontinuities in real estate and rental prices and other variables as one moves across certain locations. However, many of the neighbourhood mechanisms discussed in Durlauf (2004) and Galster (2012) are effectively externalities, and may not be well captured by real estate or rental prices.

To date, relatively little progress has been made on these two challenges. Instead, the way most empirical research defines neighborhoods seems rather ad hoc, often driven by data availability, and vary considerably across studies. The empirical work on neighborhoods and intergenerational mobility is no exception. For example, recent work based on U.S. data, which we discuss in greater detail later, uses three distinct definitions of neighborhoods. The first is the *commuting zone*, a frequent but rather coarse definition of neighborhoods. There are 741 commuting zones and, on average, each contains four counties and a population of about 380,000. This neighborhood definition might be justifiable if the focus is on geographical mechanisms of neighborhood, especially how local labor markets affect individuals' outcomes. A second definition is the *county*, which is an administrative or political subdivision of a state. County populations vary widely, with more than half the U.S. population being

38

concentrated in about 150 of the more than 3,000 counties. This definition of neighborhoods could be motivated by institutional neighborhood mechanisms, such as the availability of local public services. A third definition is the *census tract*, which provides more granularity than counties. In 2010, the U.S. was divided into about 73,000 census tracts, each with a population ranging between 1,200 and 8,000. When first delineated, census tracts were designed to be relatively homogeneous with respect to population characteristics, economic status and living conditions. Thus, using census tracts to measure neighborhood effects might be suitable if the focus is on the local social processes of the social interactive mechanism or the local physical attributes that are included in the environmental mechanisms.

Unfortunately, existing research on neighborhoods and intergenerational mobility tends to offer little if any justification other than sample sizes and statistical precision for choosing one definition of neighborhoods over another.[15] Nor does this research make clear what neighborhood mechanisms the chosen definition of neighborhoods are supposed to capture. Thus, we argue that an open and important question is how one should empirically define neighborhoods to study their impacts on intergenerational mobility.

6.2 *Moving to Opportunity Experiment*

As documented in many developed countries, children's incomes vary considerably with the area in which they grow up, even conditional on observable characteristics such as parental incomes (Chetty et al. (2014a) (U.S.); Heidrich (2017) (Sweden); Eriksen and Munk (2020) (Denmark); Acciari et al. (2019) (Italy); Deutscher and Mazumder (2020) (Australia)). These correlations could be driven by two very different sources. One possibility is that neighborhoods have causal effects on children's economic prospects. Another possibility is that the observed geographic variation is due to (unobserved) differences in the types of people living in each area.[16] To separate these two explanations, one would ideally perform an experiment that randomly assigns people to neighborhoods. While this is not feasible, it is still possible to perform an experiment that randomly incentivizes people to move to or from certain neighborhoods. One such intervention is the widely cited Moving to Opportunity (MTO) experiment, which we now discuss in greater detail.

---

[15]For example, Chetty et al. (2018, p. 1108) argue that "To maximize statistical precision, we characterize neighborhood (or "place") effects at two broad geographies: counties and commuting zones."

[16]In an attempt to control for such differences, a number of studies have used sibling comparisons within families that move across neighborhoods. This includes early work using survey data (e.g. Jencks and Mayer (1990)), and more recent work using administrative data (Chetty and Hendren (2018), Chetty et al. (2018)).

The experiment

MTO was a major randomized housing mobility experiment sponsored by the U.S. Department of Housing and Urban Development. Starting in 1994, MTO provided 4,600 low-income families with children living in public or project-based housing within some of the most disadvantaged urban neighborhoods in the U.S. the chance to move to private-market housing in much less distressed communities. A vast majority of the participating families were headed by African-American or Hispanic single mothers. The MTO Program was implemented in five large cities in the U.S. (Baltimore, Boston, Chicago, Los Angeles, and New York).

Participants in the MTO study were randomly assigned to three groups. The first was the experimental group, which received rental certificates or vouchers usable only in low-poverty areas (census tracts with less than 10 percent of the population below the poverty line in 1989). The experimental group also received counseling and assistance in finding a private unit to lease. The second group was the so-called Section 8 comparison group which received regular rental certificates or vouchers, with no conditions on where to move, and no special counseling beyond what is normally provided to voucher holders. The third was a control group that continued to receive their current project-based assistance.

What MTO Can(not) Identify?

Before we discuss the results from MTO, it is useful to make clear what the experiment can and cannot identify. We first revisit the estimating equations that are typically used in studies of MTO (see e.g. Ludwig et al. (2001); Kling et al. (2007); Ludwig et al. (2013); Chetty et al. (2016)).

The first type of estimating equation is given by

$$Y_i = \alpha_0 + \alpha_1 Exp_i + \alpha_2 S8_i + \gamma \mathbf{X}_i + \psi \mathbf{S}_i + \varepsilon_i, \tag{8}$$

where i indexes individual (or family), Y is the outcome of interest, Exp and $S8$ are indicator variables for being randomly assigned to the experimental and Section 8 groups respectively, X is a vector of baseline covariates, and S is a set of indicators for randomization site. Given the random assignment to groups, the OLS estimates of α_1 and α_2 in (8) are typically interpreted as the causal impacts of being offered a given type of voucher to move through MTO.

Interpreting the magnitude of the estimates of (8) is difficult because not all families

offered vouchers actually took them up. To address this issue, the MTO studies typically instrument for voucher take-up with the group assignment indicators. Formally, they estimate specifications of the form

$$y_i = \beta_0 + \beta_1 \text{Take Exp}_i + \beta_2 \text{ Take } S8_i + \theta \mathbf{X}_i + \delta \mathbf{S}_i + \varepsilon_i, \tag{9}$$

where *TakeExp* and *TakeS*8 are indicators for taking up the experimental and Section 8 vouchers, respectively. To address that these are endogenous variables, (9) is estimated using two-stage least squares with the randomly-assigned group assignment indicators *Exp* and *S*8 as instruments for taking up the vouchers.

A causal interpretation of the two-stage least squares estimates requires several strong assumptions. First, it is necessary to assume that the MTO voucher offers only affect outcomes through the actual use of the voucher to lease a new residence. However, families in the experimental group had additional access to life-skills counseling sessions that could directly affect their outcomes. Another issue is that causal interpretation of two-stage least squares estimates in settings with multiple endogenous variables requires strong assumptions over and above those needed in settings with a binary treatment variable. This issue is ignored in the existing research on MTO, which erroneously apply the logic of two-stage least squares with a binary treatment variable to the MTO setting with two treatment variables. However, as shown in Kirkeboen et al. (2016), even with a valid instrument for both treatment variables a causal interpretation of the 2SLS estimates requires either information about or strong restrictions on individuals' preferences, or an assumption of constant effects of each treatment variable across individuals.

Key Empirical Findings from MTO

A number of studies have used specifications of the form (8) and (9) to analyze the MTO experiment (see e.g. Ludwig et al. (2001); Kling et al. (2007); Ludwig et al. (2013); Chetty et al. (2016)). Instead of going into the details of each study, we highlight a few key insights about how MTO affected voucher-take up, the probability of moving, and the choice of neighborhood, as well the short and longer run outcomes of children and adults.

One insight from the evaluations of MTO (see e.g. Ludwig (2012)) is that even though imperfect compliance to treatment assignment is empirically important, the experiment did succeed in changing the neighborhood characteristics of the participants. Only around 47% of families who were offered an experimental group voucher and 63% of those offered a Section 8 group voucher relocated through MTO. Yet, the MTO induced moves caused significant and

meaningful differences in neighborhood characteristics between the experimental, Section 8, and control groups. In the short-run, families in the experimental group and families in the Section 8 group were, on average, living in census tracts with poverty rates that were 17 and 14 percentage points lower than families in the control group. The neighborhoods of these families also differed in several other ways as a result of the MTO experiment, such as with respect to crime and social environment. Over time, however, the differences across the groups change and to some extent attenuate. This is in part because of large underlying mobility among the participants in the MTO experiment, but also due to significant changes to and improvements in the neighborhoods of the control group families.

Another set of findings from the MTO studies is that assignment to the experimental group or the Section 8 group has significant and positive effects on children but little if any measurable impact on parents.[17] While most of this research looks at impacts in the short and medium run, Chetty et al. (2016) consider the long term effects. Their findings suggest that young children assigned to the experimental group experience meaningful increases in educational attainment and earnings as adults compared to young children in the control group.

Finally, some of the findings from the MTO are at odds with the conclusion in previous observational studies of significant effects of exposure to poor neighborhoods (see e.g. Elliott (1999); Fauth et al. (2004); Shang (2014); Vartanian (1997)). Harding et al. (2021) examine whether it is possible to reconcile these findings by replicating the MTO experimental estimate by applying non-experimental methods to both the PSID data and the MTO data. They conclude that hypotheses related to effect heterogeneity, treatment magnitude, treatment duration, neighborhood effect nonlinearities, residential mobility, and neighborhood measures are unlikely to account for differences between experimental and non-experimental estimates. Instead, they argue that the differences are likely to be due to selection bias in the non-experimental studies. However, the statistical power of the analyses severely limits the confidence one can have in these conclusions, as recognized by Harding et al. (2021).

When interpreting the findings from MTO, it is important to keep in mind what this experiment can and cannot identify. In contrast to what is often claimed, comparisons of outcomes across the MTO groups are unlikely to uncover the effects of particular neighborhood characteristics, such as low poverty.[18] The comparisons combine exposure effects from moving

---

[17]Some studies do report statistically significant effects for a few parental outcomes (see e.g. Ludwig et al. (2013)). However, multiple testing is a serious concern given the large number of outcomes that have been considered both within and across various studies.

[18]A notable example of such claims can be found at the NBER's website for the MTO experiment: "Because of the random assignment design, the MTO study generates comparable groups of adults and children living

to a particular type of neighborhood with counseling and assistance as well as any disruption costs of moving to a given neighborhood. Furthermore, many families in both the control and treatment groups moved frequently across neighborhoods after the experiment, making it very difficult to isolate the effects of the time spent in a particular type of neighborhood. In addition, the moves induced by MTO change an entire bundle of neighborhood characteristics. Thus, it is not possible to disentangle changes in neighborhood poverty from simultaneous changes in other characteristics that could influence individual's outcomes. For example, the reductions in asthma rates may be due to improvements in housing quality (asthma is strongly associated with rat infestations) and nothing about the neighborhood per se (Durlauf, 2004).

For these reasons, one might argue that the MTO experiment is primarily informative about the socio-economic consequences of being offered housing vouchers, not neighborhood effects. However, it is important to keep in mind that MTO changes the type—but not necessarily the amount—of housing support an individual may receive. Thus, MTO does not inform us about the consequences of expanding a voucher program to include currently ineligible families. Instead, the policy question that the MTO results might be closest to answering is what would happen if we changed the mix of means-tested housing programs to include a larger share of housing vouchers and a smaller share of project-based units.

6.3 Upward Mobility of Individual Neighborhoods

The purview of much of the empirical research on neighborhood effects is limited to estimating *average* impacts of living in or moving to a given type of neighborhood (e.g. low poverty).[19] Over the past few years, however, we have seen a very ambitious and highly influential body of empirical work that tries to estimate how *individual* neighborhoods correlate with and possibly affect children's outcomes. This literature started with the work of Chetty et al. (2014a), Chetty et al. (2018) and Chetty and Hendren (2018), who asked the question: Where in the United States is the land of opportunity?

To answer this question, Chetty et al. (2014a), Chetty et al. (2018) and Chetty and Hendren (2018) use large administrative data to analyze how intergenerational income mobility

---

in different types of neighborhoods, so that a comparison of outcomes across research groups can uncover the potential effects of neighborhood characteristics across a range of family and children's outcomes." See Pinto (2019) for an econometric analysis of MTO that is explicit on the assumptions needed to draw causal inference about effects of residing in different neighborhoods types.

[19]There is also a literature that seeks to quantify how much of the variation in children's outcomes can be explained by neighborhood effects. To this end, these studies compare the covariance of outcomes for siblings in the same community and a pair of unrelated individuals in the same community. See e.g. Solon et al. (2000) for evidence from the U.S. and Raaum et al. (2006) for evidence from Norway. Durlauf (2004) offers a critical review of this literature.

vary across areas in the United States. The key outputs from these studies were "local statistics" on upward mobility across commuting zones, counties, and census tracts.[20] The stated goals were to draw the attention of policymakers to specific low-mobility neighborhoods in the United States that need improvement and to help low-income families move to specific high-mobility neighborhoods.[21] Below, we investigate the extent to which these goals were achieved by critically assessing the credibility and informativeness of the estimated local statistics on upward mobility.

Estimates of Individual Neighborhood Effects

Broadly speaking, the work of Chetty and coauthors uses federal income tax records spanning 1996–2012 to construct two types of estimates. The first type is the correlational estimates, as reported in Chetty et al. (2014a), Chetty and Hendren (2018) and Chetty et al. (2018). In this work, the authors document how children's expected incomes conditional on their parents' incomes vary according to the area (commuting zone (CZ), county, or census tract) in which they grew up. The second is the mover estimates, as reported in Chetty and Hendren (2018) and Chetty et al. (2018). The goal of this work is to draw causal inference about the effects of neighborhoods on children's outcomes. We now present the two type of estimates in turn.

Correlational Estimates of Neighbourhood Effects.   In the baseline analysis, Chetty et al. (2014a, 2018) define the following measure of intergenerational mobility:

$$\bar{y}_{cp} \equiv E\left[y_i \mid c(i) = c, p(i) = p\right],$$

where y_i is child i's percentile rank in the national distribution of incomes relative to all others in her birth cohort; child i's income is measured as her average income in the years 2014–2015 (aged 31–37 depending on cohort); $p(i)$ denotes the child's parental income percentile in the national distribution of parental income in child i's birth cohort; and $c(i)$ is the area in which child i grew up. They focus on $\bar{y}_{c25}$, the expected income rank of children who grew up in area c with parents at the 25th percentile of the national income distribution of parental

---

[20]The chief novelty of this work is arguably the granular estimates of intergenerational mobility. For a description of intergenerational mobility across broader regions of the U.S., see for example Connor and Storper (2020) and the references therein.

[21]Recently, local statistics on upward mobility have also been produced in many other developed countries, including Sweden, Denmark, Australia and Italy (Heidrich, 2017; Eriksen and Munk, 2020; Deutscher and Mazumder, 2020; Acciari et al., 2019)). These studies usually find smaller, but still significant regional differences, while Acciari et al. (2019) show a strong North-South gradient in Italy.

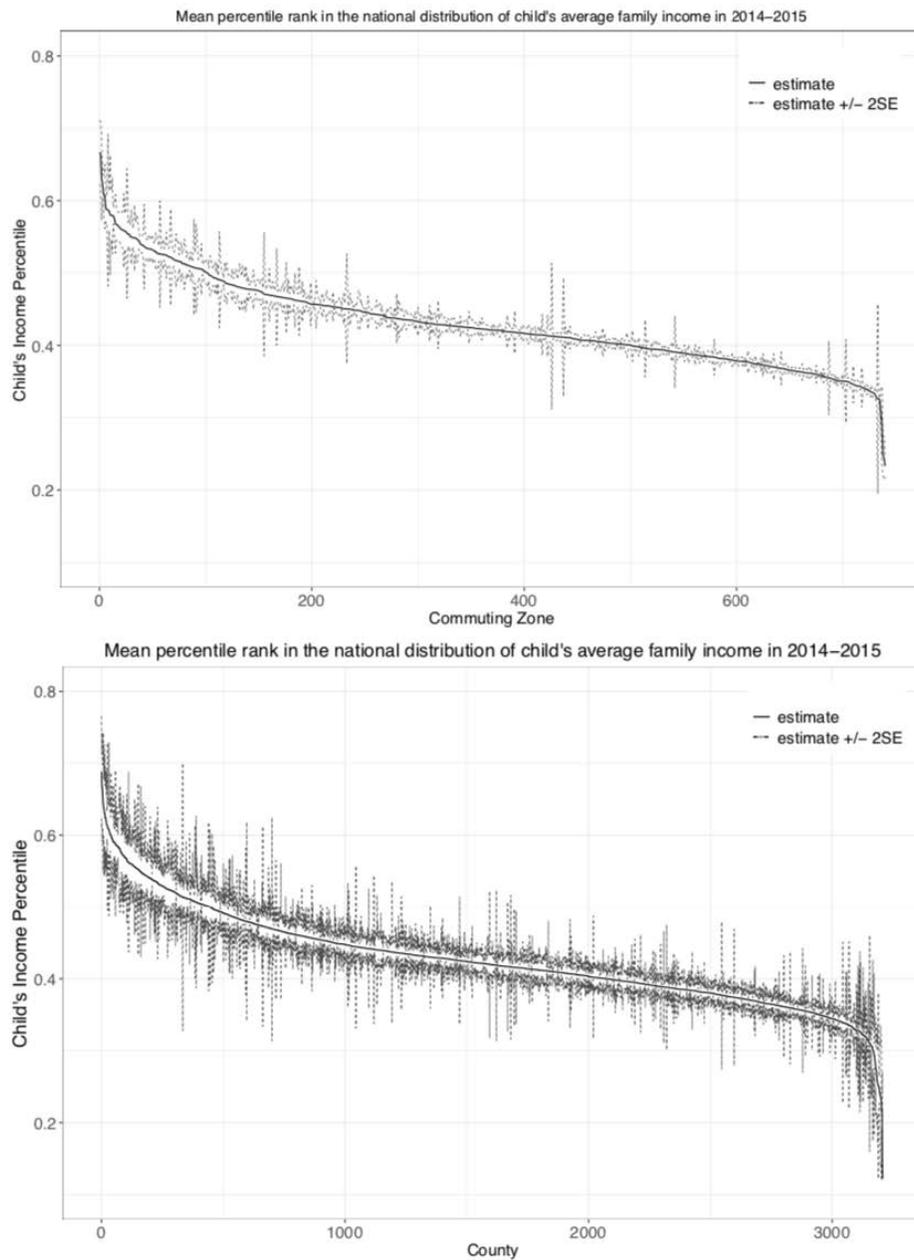income. We refer to Chetty et al. (2014a, 2018) for estimation details.

Figure 3 presents the estimates of $\bar{y}_{c25}$ with marginal confidence intervals (estimates plus or minus twice the standard errors), as reported in Chetty et al. (2018). These correlational estimates of upward mobility cover all 741 commuting zones and 3208 of the 3219 counties. Places are sorted by their values of $\hat{\bar{y}}_{c25}$, and then point estimates and their marginal confidence intervals for each CZ (top graph) and county (bottom graph) are reported. There is considerable variation in $\hat{\bar{y}}_{c25}$ across areas. Since CZs typically comprise several counties, it is not surprising that the standard errors tend to be a lot larger when a neighborhood is defined as a county rather than as a CZ.

To analyze whether an individual CZ (or county) has high or low mobility compared to other CZs (or counties), Chetty et al. present heat maps based on the point estimates of upward mobility. They construct these maps by dividing the CZs (or the counties) into deciles based on their estimated value of $\bar{y}_{c25}$. Figure 4 presents the heat map for the CZs. This map is the same as presented in Chetty et al. (2014a). Lighter colors represent deciles with higher values of $\bar{y}_{c25}$. Equivalently, one can interpret the heatmap as showing the ranks of CZs by assigning the same color to ranks in a decile to easy readability (rather than a unique color to each rank).

As is evident from Figure 4, the point estimates of upward mobility vary significantly across areas, even within a state. For example, CZs in the top decile have a $\hat{y}_{c25} > 0.517$, while those in the bottom decile have $\hat{y}_{c25} < 0.362$. Note that the 36th percentile of the family income distribution for children at age 31–37 is $26,800$, while the 52nd percentile is $44,800$; hence, the differences in upward mobility across these areas correspond to substantial differences in children's incomes. Comparisons such as this lead Chetty et al. (2014a, p. 1554) to conclude that "intergenerational mobility varies substantially across areas within the United States. For example, the probability that a child reaches the top quintile of the national income distribution starting from a family in the bottom quintile is 4.4 percent in Charlotte but 12.9 percent in San Jose."[22]

In subsequent work, Chetty et al. (2018) shift attention from individual CZs and counties to individual census tracts. For each tract, they estimate children's earnings distributions, incarceration rates, and other outcomes in adulthood by parental income, race, and gender. These estimates, they argue, allow them to trace the roots of outcomes such as poverty and

---

[22]Chetty et al. also explore the factors correlated with upward mobility, finding that high mobility CZs have less residential segregation, less income inequality, better primary schools, greater social capital, and greater family stability. However, these are simply correlates of upward mobility and difficult to interpret either causally and economically.

Mean percentile rank in the national distribution of child's average family income in 2014–2015

Mean percentile rank in the national distribution of child's average family income in 2014–2015

Note: The estimates cover all 741 commuting zones (Top) and 3208 of the 3219 counties (Bottom). Source: Mogstad et al. (2020).

Figure 3. Estimates of $\bar{y}_{c25}$, the expected percentile rank of a child's average household income in the national distribution of her cohort, with marginal confidence intervals (estimates plus or minus twice the standard errors)
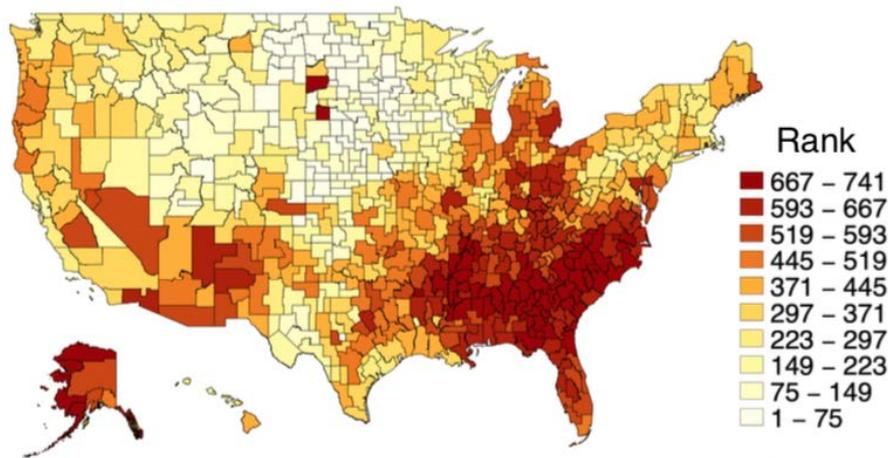
Note: The heat map is based on estimates of $\bar{y}_{c25}$, the mean percentile rank of child's average household income for 2014–2015, for the full set of CZs. The map is constructed by dividing the CZs into deciles based on the estimated values of $\bar{y}_{c25}$, and shading the areas so that lighter colors correspond to higher absolute mobility or, equivalently, lower ("better") rank. Source: Mogstad et al. (2020).

Figure 4. Ranking of Commuting Zones by point estimates and lower and upper endpoints of simultaneous confidence sets.

incarceration back to the neighborhoods in which children grew up. They find that the point estimates of upward mobility vary greatly even across nearby tracts, as illustrated by the heat map of Brooklyn's Brownsville neighborhood in Figure 5. Based on these comparisons, Chetty et al. (2018) conclude that neighborhoods at a very granular level may be very important for children's outcomes.

Mover Estimates of Neighbourhood Effects. Although Chetty et al. (2014a, 2018) on occasion use a language of cause and effect to describe their findings, they are careful to emphasize that the correlational estimates of upward mobility across areas cannot necessarily be given a causal interpretation. To address concerns about selection bias, Chetty and Hendren (2018) study how the outcomes of children who move across CZs vary with the age at which they move. The parameters of interest are the exposure effects of spending an additional year of one's childhood in a given area. The identifying assumption is that the selection effects associated with the family and child are independent of child's age when the family moves.[23]

---

[23] Chetty and Hendren (2018) take several steps to examine and relax this assumption, such as including family fixed effects so that one compares outcomes across siblings who differ in their exposure to neighborhoods. Heckman and Landersø (2020) critically examine identification and data issues in the work of Chetty et al.

Household Income
for Black Children of Low Income Parents
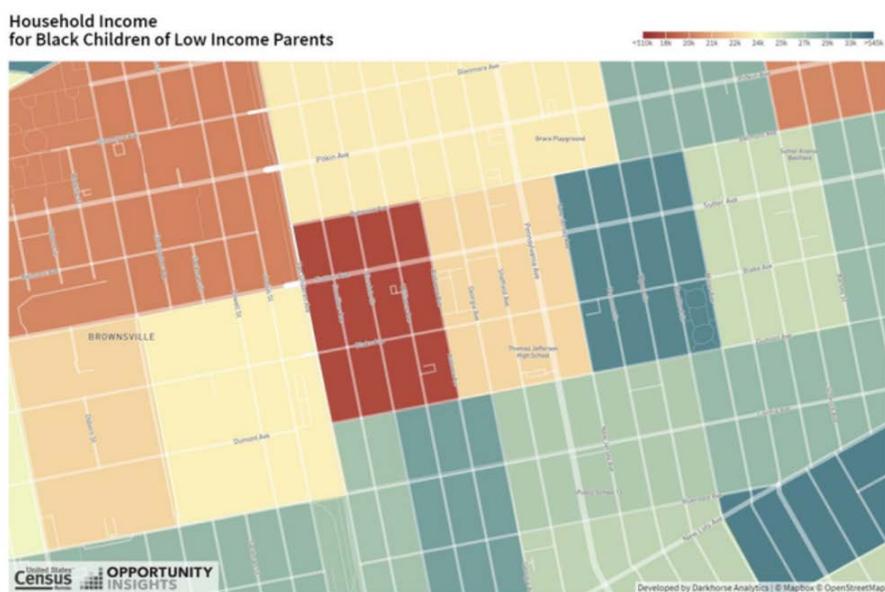
Figure 5. Disparities in average earnings for people born into low-income black families and raised in different, but nearby, parts of Brooklyn's Brownsville neighborhood where each rectangle represents a city block.

To properly define this parameter and show how it is recovered from the data, some notation is useful. Consider a child i from a set of one-time movers from an origin $o(i)$ to a destination $d(i)$. She moves at the age $m(i)$ and spends $A - m(i)$ time in the destination. The (vector of the) amount of time spent in a given area is denoted by

$$e_{ic} \equiv \begin{cases} A - m_i & c = d(i) \\ m_i & c = o(i) \\ 0 & \text{elsewhere.} \end{cases}$$

The exposure effects can be estimated by the regression model

$$y_i = \alpha_{od} + \vec{e}_i \cdot \vec{\mu} + \varepsilon_i,$$

(2014a, 2018) and Chetty and Hendren (2018). See also Durlauf (2004) for a discussion of challenges to using sibling comparisons and movers to draw inference about neighborhoods effects.

where α_{od} is an origin-by-destination fixed effect, $\overrightarrow{e}_i \equiv (e_{ic} : c = 1, 2, \ldots)$ is a vector of explanatory variables for the number of years that child i lived in place c during her childhood, and the exposure effects are given by the parameters $\overrightarrow{\mu} \equiv (\mu_{cp} : c = 1, 2, \ldots) \equiv (\mu_c^0 + \mu_c^1 p : c = 1, 2, \ldots)$, where p is the parental income percentile. The estimates are normalized to be mean zero across places, so that μ_{cp} measures the exposure effect relative to the average place. Chetty and Hendren (2018) focus on μ_{c25}, the effect of spending an additional year of childhood in area c for children with parents at the 25th percentile of the national income distribution of parental income.
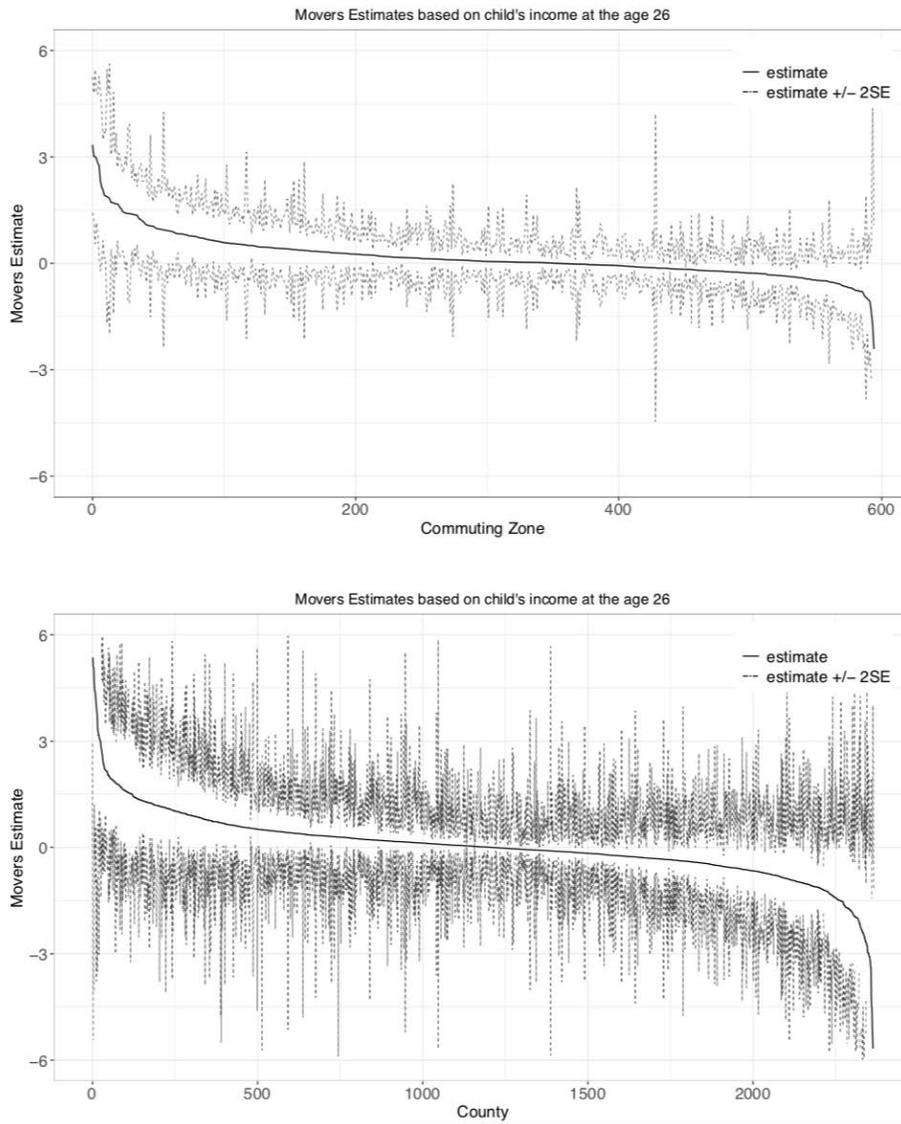
Figure 6 presents the point estimates of exposure effects μ_{c25} with marginal confidence intervals (estimates plus or minus twice the standard errors) from Chetty and Hendren (2018). These results cover 595 of the 741 CZs and 2367 of the 3219 counties.[24] The point estimates suggest a lot of variation in exposure effects across areas. However, the standard errors are large, especially at the county level. Recognizing the large statistical uncertainty, Chetty and Hendren (2018) construct forecasts that minimize the mean-squared-error of the predicted impact of growing up in a given place. The stated purpose of these forecasts is to guide families seeking to move to better neighborhoods.

It is important to observe, however, that these forecasts are very similar to the correlational estimates in most CZs and counties. The forecasts are constructed as weighted averages of the correlational estimates (based on stayers) and the mover estimates, with greater weight on the mover estimates when they are more precisely estimated. Given that most mover estimates of μ_{c25} are very noisy, the forecast estimates should be very similar to the correlational estimates. Indeed, Mogstad et al. (2020) calculate that in a majority of the CZs, the forecasts assign at least 90 percent of the weight to the correlational estimates. Unfortunately, Chetty and Hendren (2018) do not report confidence intervals on the forecast estimates, and, thus, there is no telling how precise these forecasts are.

Do the neighbourhood estimates reflect noise or signal?

A much celebrated conclusion of Chetty et al. (2014a, 2018) and Chetty and Hendren (2018) is that the choice of local neighborhood can have a large impact on children's outcomes. However, this conclusion does not take into account the large statistical uncertainty surrounding

---

[24]Chetty and Hendren (2018) choose not to report results for the other counties and CZs due to limited data in these areas. However, they estimate exposure effects at the census tract level (for origin-to-destinations which have at least 20 observations) in Chetty et al. (2018). For brevity, we chose not to discuss these results, as the estimates are very noisy. Instead, we focus on the CZs and counties for which informative conclusions are more likely.

Note: The estimates cover 595 of the 741 commuting zones (Top) and 2367 of the 3219 counties (Bottom).
Source: Mogstad et al. (2020)

Figure 6. Movers estimates of exposure effects μ_{c25} with marginal confidence intervals (estimates plus or minus twice the standard errors).

the neighbourhood estimates. This uncertainty raises the question of how informative these local statistics are about a given neighborhood having relatively high or low income mobility compared to other neighborhoods.

Two recent papers develop and apply new methods to investigate this question. The first is by Mogstad et al. (2020), who develop a new method to account for statistical uncertainty in rankings. The second paper is by Andrews et al. (2020), showing how to conduct inference about a parameter that is chosen as the "best" out of a set of choices, where the choice must be based on point estimates as opposed to the true values. Mogstad et al. (2020) apply both these methods to re-examine the findings of Chetty et al. (2014a, 2018) and Chetty and Hendren (2018). Mogstad et al. (2020) conclude that many (but not all) of the celebrated findings about neighborhoods and intergenerational mobility are not robust to taking uncertainty into account. We next show how they arrive at this conclusion.

Ranking of Neighborhoods by Upward Mobility. Mogstad et al. (2020) apply their procedure to the point estimates and standard errors from Chetty et al. (2018) and Chetty and Hendren (2018), reported in Figures 3 and 6. They compute (i) the marginal confidence sets for the rank of a given place, (ii) the simultaneous confidence sets for the ranks of all places, and (iii) the confidence sets for the τ-best (or the τ-worst) ranked places, emphasizing that (i)–(iii) answer distinct economic questions. Marginal confidence sets answer the question of whether a given place has relatively high or low upward mobility compared to other places. Thus, (i) is relevant if one is interested in whether a particular place is among the worst or the best places to grow up in terms of upward mobility. Simultaneous confidence sets allow such inferences to be drawn simultaneously across all places. Thus, (ii) is relevant if one is interested in broader geographic patterns of upward mobility across the United States. By comparison, confidence sets for the τ-best (or τ-worst) answer the more specific question of which places cannot be ruled out as being among the areas with the most (least) upward mobility. In other words, (iii) is relevant if one is interested in only the top (or bottom) of a league table of neighborhoods by upward mobility.

The analyses of Mogstad et al. (2020) reveal that the most robust findings are obtained if one restricts attention to the 50 most populous CZs or counties (as in parts of the analyses of Chetty et al. (2014a)). In that case, both the marginal and simultaneous confidence sets are relatively narrow, and few places cannot be ruled out as being among the top or bottom five. By comparison, in the national ranking of all commuting zones or counties by upward mobility, it is often not possible to determine with statistical confidence whether a given place has relatively high or low income mobility compared to other places. Thus, it is not possible
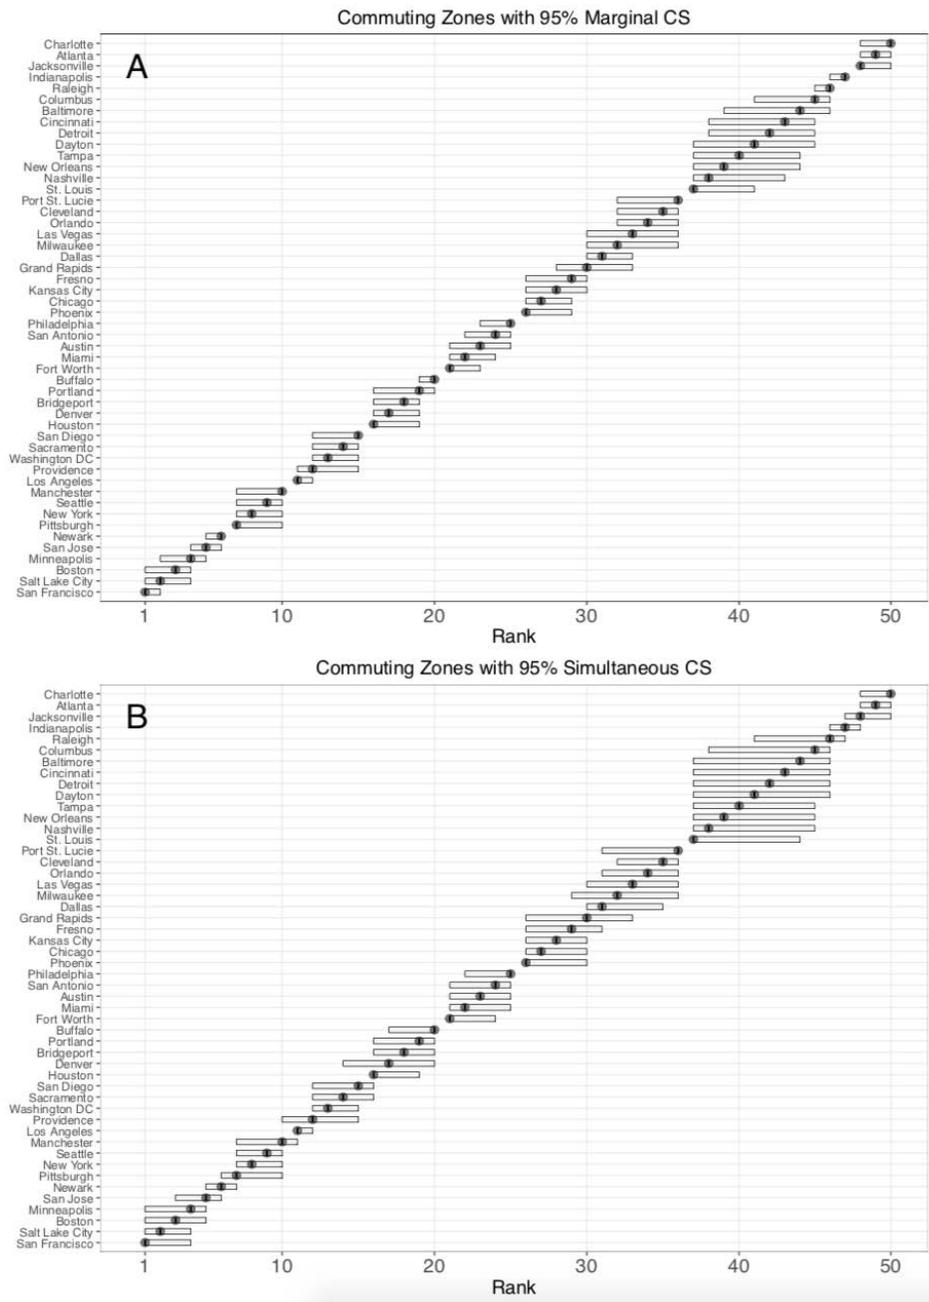
Note: **Panel A**: point estimates and the 95% marginal confidence sets ("CS") for the ranking of the 50 most populous CZs by $\bar{y}_{c25}$. **Panel B**: point estimates and the 95% simultaneous confidence sets ("CS") for the ranking of the 50 most populous CZs by $\bar{y}_{c25}$. Source: Mogstad et al. (2020)

Figure 7. Ranking of CZs according to the point estimates of $\bar{y}_{c25}$.

Table 2. Ranking Results for the Top and Bottom Five CZs

Panel A: Top 5

		Correlational					Movers			
Rank	τ	CZ	$\hat{\bar{y}}_{c25}$	SE	95% CS	τ-best	CZ	$\hat{\mu}_{c25}$	SE	95% CS	τ-best
1	1	San Francisco	0.457	0.001	[1, 2]	4	Seattle	0.229	0.082	[1, 38]	44
2	2	Salt Lake City	0.457	0.001	[1, 4]	4	Washington DC	0.163	0.077	[1, 41]	48
3	3	Boston	0.453	0.001	[1, 4]	5	Cleveland	0.124	0.107	[1, 48]	50
4	4	Minneapolis	0.452	0.001	[2, 5]	5	Fort Worth	0.121	0.090	[1, 48]	50
5	5	San Jose	0.449	0.001	[4, 6]	6	Minneapolis	0.116	0.120	[1, 48]	50

Panel B: Bottom 5

		Correlational					Movers			
Rank	τ	CZ	$\hat{\bar{y}}_{c25}$	SE	95% CS	τ-worst	CZ	$\hat{\mu}_{c25}$	SE	95% CS	τ-worst
46	5	Raleigh	0.369	0.001	[45, 46]	10	Charlotte	-0.248	0.096	[3, 50]	49
47	4	Indianapolis	0.364	0.001	[46, 47]	5	Port St. Lucie	-0.263	0.090	[3, 50]	49
48	3	Jacksonville	0.358	0.001	[48, 50]	4	Raleigh	-0.278	0.105	[3, 50]	49
49	2	Atlanta	0.358	0.001	[48, 50]	3	Fresno	-0.377	0.100	[13, 50]	48
50	1	Charlotte	0.355	0.001	[48, 50]	3	New Orleans	-0.391	0.111	[14, 50]	48

Note: **Panel A**: Top 5 among the 50 most populous commuting zones ranked by the correlational estimates on the left and by the movers estimates on the right. **Panel B**: Bottom 5 among the 50 most populous commuting zones ranked by the correlational estimates on the left and by the movers estimates on the right. "95% CS" refers to the 95% marginal confidence set for the rank, and "τ-best" and "τ-worst" refer to the size of the 95% confidence sets for the "τ-best" and "τ-worst" commuting zones. Source: Mogstad et al. (2020) .

to give an informative answer to the question of where in the United States is the land of opportunity. Another key finding is that the rankings of even the most populous commuting zones or counties become largely uninformative if one uses movers across areas to address concerns about selection bias and draw causal conclusions.

To illustrate how Mogstad et al. (2020) produce these findings, it is useful to revisit their analysis of the 50 most populous CZs. Figure 7 presents the ranking of these CZs according to the point estimates of $\bar{y}_{c25}$. Panel A displays the marginal confidence sets while Panel B reports the simultaneous confidence sets. Table 2 reports additional results for the top five CZs (Panel A) and the bottom five CZs (Panel B). Each panel of this table presents two sets of results: Columns 3–7 are based on the correlational estimates of upward mobility $\bar{y}_{c25}$, while columns 8–12 are based on the movers estimates of exposure effects μ_{c25}. For each set of results, they report the point estimates, the standard errors, the 95% marginal confidence sets, and the number of places in the 95% confidence sets for the τ-best (top panel) or the τ-worst values of $\bar{y}_{c25}$ or μ_{c25}.

Among the 50 largest CZs by population size, the point estimates of $\bar{y}_{c25}$ range from 0.457 in San Francisco to 0.355 in Charlotte. As evident from Panel A of Figure 7, the marginal

confidence sets based on the correlational estimates are relatively narrow, especially for the CZs at the top and the bottom of the ranking. This finding suggests that citizens of these CZs can be quite confident in the mobility ranking of their hometown. For instance, with 95% confidence, San Francisco is among the top two of these 50 CZs in terms of income mobility. By comparison, with 95% confidence, Charlotte is among the bottom three of these 50 CZs in terms of income mobility.

A natural question is whether the ranking of the CZs according to the correlational estimates remains informative if one allows inferences to be drawn simultaneously across the 50 CZs. The results in Panel B of Figure 7 suggest this is indeed the case and one can have high confidence about which CZs are at the top and bottom of the correlational ranking. The sizes of the 95% confidence sets for the τ-best and τ-worst CZs confirm this finding. For example, only four (three) places cannot be ruled out as being among the top (bottom) two CZs in terms of income mobility. Furthermore, there are only six places that cannot be ruled out as being among the top five CZs, while ten CZs cannot be ruled out as being among the bottom five places.

Taken together, the results based on the correlational estimates $\bar{y}_{c25}$ suggest it is possible to achieve a quite informative ranking of the 50 largest CZs according to upward mobility. In contrast, the exposure effects μ_{c25} are too imprecisely estimated to draw firm conclusions about which CZs produce more or less upward mobility. As evident from the marginal confidence sets for μ_{c25} in column 11 of Table 2, it is difficult to learn much about whether a particular CZ has relatively high or low exposure effects. For example, the citizens of Seattle cannot rule out with 95% confidence that the majority of other CZs have higher income mobility. Drawing inferences simultaneously across all CZs is even more challenging, as evident by the τ-best and τ-worst results for μ_{c25}. Consider, for example, column 12 of Panel A in Table 2. As these results show, none of the 50 CZs can be ruled out with 95% confidence as being among the top five places in terms of exposure effects.

Inference on neighborhoods with the highest estimated upward mobility. Mogstad et al. (2020)'s methods and those developed by Andrews et al. (2020) share some technical similarities, but answer distinct economic questions, and should thus be viewed as complements, not substitutes. To illustrate this empirically, Mogstad et al. (2020) apply the methods of Andrews et al. (2020) to the correlational estimates and construct 95% confidence sets for the true mobility of the CZ with the highest estimated mobility. For instance, their 95% confidence set on the expected income rank of children with parents at the 25 percentile who grew up in the "winning" CZ (among all CZs) is (0, .66). Since the confidence set includes zero

(the smallest possible value of the mobility measure), one cannot be confident that the CZ with the highest point estimate truly has high mobility. The corresponding confidence set for the "winning" CZ among only the 50 most populous CZs, San Francisco, is $(0.389, 0.457)$. While this confidence set excludes zero, comparing it to the range of estimates for the 50 most populous CZs, it is still fairly wide.

Taken together, the results using Andrews et al. (2020)'s methods suggest there is considerable statistical uncertainty about the true value of upward mobility at the top of the estimated ranking of CZs, even if one restricts the study to the 50 most populous CZs. This conclusion holds if one considers neighborhoods at a more granular level (counties, Census tracts), or if one uses movers across areas to address concerns about selection.

Discussion of Policy Implications of Neighborhood Estimates

The neighborhood estimates of Chetty et al. (2014a, 2018) and Chetty and Hendren (2018) have been highly influential both among policymakers and among researchers. For example, the rankings of neighborhoods by (point estimates of) intergenerational mobility play a key role in Chetty's 2014 Testimony for the United States Senate Committee on the Budget (Chetty, 2014). In this testimony (pages 6 and 7), he emphasizes that policy should target areas that are ranked at the bottom of the league tables based on their estimates of upward mobility:

> "Since rates of upward mobility vary widely across cities, place-based policies that focus on specific cities such as Charlotte or Milwaukee may be more effective than addressing the problem at a national level."

Moreover, Chetty claims that it is key to disseminate information about which areas have relatively high and low estimates of upward mobility:

> "Perhaps the most cost-effective way to improve mobility may be to publicize local statistics on economic mobility and other related outcomes. Simply drawing attention to the areas that need improvement can motivate local policy makers to take action. Moreover, without such information, it is difficult to determine which programs work and which do not. The federal government is well positioned to construct such statistics at minimal cost with existing data. The government could go further by offering awards or grants to areas that have substantially improved their rates of upward mobility. Shining a spotlight on the communities where children have opportunities to succeed can enable others to learn

55

from their example and increase opportunities for economic mobility throughout America."

In light of the large degree of uncertainty, however, one may be concerned that local statistics do not necessarily contain valuable information about upward mobility. As a consequences of this uncertainty, it can be problematic to use such statistics to disseminate information or target interventions. The spotlight might be shining on noise, not signal.

To illustrate this point, Mogstad et al. (2020) revisit the recent Creating Moves to Opportunity Experiment (CMTO) of Bergman et al. (2019). With the aim of helping families move to neighborhoods with higher mobility rates, the authors conduct a randomized controlled trial with housing voucher recipients in Seattle and King County. A treatment group of low-income families were offered assistance and financial support to find and lease units in areas that were classified as high upward-mobility neighborhoods within the county. Bergman et al. (2019) define high upward-mobility neighborhoods as Census tracts with point estimates of upward mobility among the top one-third of the tracts in the county, ignoring the large statistical uncertainty of the estimates. Since no data on outcomes is yet available, the authors predict the impacts of the moves induced by the CMTO program on children's future outcomes using the point estimates of upward mobility of the individual tracts. However, Mogstad et al. (2020) show that the classification of a given tract as a high upward-mobility neighborhood may simply reflect statistical uncertainty, not that mobility is particularly high in that neighborhood.[25] As discussed in greater detail in their paper, this noise in the classification raises the question of whether one could be confident that CMTO would actually help families move to high upward-mobility neighborhoods, prior to the experiment taking effect.

7 Conclusion

The human capital approach in economics considers how the productivity of people in market and non-market situations is changed by investments in education, skills, and knowledge. Half a century ago, Becker and Tomes (1979, 1986) developed the human capital approach into a general theory for income inequality, both across families within a generation and between different generations of the same family. Much of the progress since has focused on improving measurements, uncovering new facts, or identifying causal impacts of various

---

[25]In addition to the problem of uncertainty, one may worry about using correlational estimates to define a treatment group and that the historical estimates do not capture the current mobility of a census tract.

determinants of mobility. The goal of our article was to critically review some of the contributions that have been made on these fronts. In the course of doing so, we also highlighted some of the limitations or weaknesses with the current empirical literature on intergenerational mobility. We believe important progress can be made by combining theory, econometrics and empirics in a coherent and transparent way. Doing so may get us closer to fulfilling the goal of Becker and Tomes (1986, p. 3) of an *"analysis that is adequate to cope with the many aspects of the rise and fall of families"*.

References

Abbiati, G. and Barone, C. (2017). Is university education worth the investment? The expectations of upper secondary school seniors and the role of family background. *Rationality and society*, 29(2):113–159.

Acciari, P., Polo, A., and Violante, G. (2019). "And yet it moves": Intergenerational mobility in Italy. Technical Report 25732, National Bureau of Economic Research.

Agostinelli, F. and Sorrenti, G. (2018). Money vs. time: family income, maternal labor supply, and child development. Technical Report 273, University of Zurich, Department of Economics, Working Paper.

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.

Almond, D., Currie, J., and Duque, V. (2018). Childhood circumstances and adult outcomes: Act II. *Journal of Economic Literature*, 56(4):1360–1446.

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. Technical report, IZA Institute for Labor Economics.

Andrews, I., Kitagawa, T., and McCloskey, A. (2020). Inference on winners. Technical Report 25456, National Bureau of Economic Research.

Attanasio, O. and Kaufmann, K. (2009). Educational choices, subjective expectations, and credit constraints. Technical report, National Bureau of Economic Research.

Barth, D., Papageorge, N. W., and Thom, K. (2020). Genetic endowments and wealth inequality. *Journal of Political Economy*, 128(4):1474–1522.

Becker, G. S. (1959). Union restrictions on entry. *The public stake in union power*, pages 209–24.

Becker, G. S. (1981). *A Treatise on the Family*. Harvard University Press.

Becker, G. S., Kominers, S. D., Murphy, K. M., and Spenkuch, J. L. (2018). A theory of intergenerational mobility. *Journal of Political Economy*, 126(S1):S7–S25.

Becker, G. S. and Tomes, N. (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of political Economy*, 87(6):1153–1189.

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.

Behrman, J. R. and Rosenzweig, M. R. (2002). Does increasing women's schooling raise the schooling of the next generation? *American economic review*, 92(1):323–334.

Behrman, J. R. and Taubman, P. (1989). Is schooling "mostly in the genes"? Nature-nurture decomposition using data on relatives. *Journal of Political economy*, 97(6):1425–1446.

Benjamin, D. J., Cesarini, D., Chabris, C. F., Glaeser, E. L., Laibson, D. I., Age, G. S.-R. S., Guðnason, V., Harris, T. B., Launer, L. J., Purcell, S., et al. (2012). The promises and pitfalls of genoeconomics. *Annu. Rev. Econ.*, 4(1):627–662.

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L., and Palmer, C. (2019). Creating Moves to Opportunity: Experimental evidence on barriers to neighborhood choice. Technical Report 26164, National Bureau of Economic Research.

Björklund, A. and Jäntti, M. (1997). Intergenerational income mobility in Sweden compared to the United States. *The American Economic Review*, 87(5):1009–1018.

Björklund, A. and Jäntti, M. (2012). How important is family background for labor-economic outcomes? *Labour Economics*, 19(4):465–474.

Björklund, A. and Jäntti, M. (2020). Intergenerational mobility, intergenerational effects, sibling correlations, and equality of opportunity: a comparison of four approaches. *Research in Social Stratification and Mobility*, page 100455.

Björklund, A., Jäntti, M., et al. (2009a). Intergenerational income mobility and the role of family background. *Oxford handbook of economic inequality*, 491:521.

Björklund, A., Jäntti, M., and Lindquist, M. J. (2009b). Family background and income during the rise of the welfare state: brother correlations in income for Swedish men born 1932–1968. *Journal of Public Economics*, 93(5-6):671–680.

Bjorklund, A., Jantti, M., and Solon, G. (2005). Influences of nature and nurture on earnings variation. *Unequal chances: Family background and economic success*, pages 145–164.

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.

Björklund, A. and Salvanes, K. G. (2011). Education and family background: Mechanisms and policies. In *Handbook of the Economics of Education*, volume 3, pages 201–247. Elsevier.

Black, S. E., Devereux, P. J., Lundborg, P., and Majlesi, K. (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.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005). Why the apple doesn't fall far: Understanding intergenerational transmission of human capital. *American economic review*, 95(1):437–449.

Blanden, J. (2013). Cross-country rankings in intergenerational mobility: a comparison of approaches from economics and sociology. *Journal of Economic Surveys*, 27(1):38–73.

Bleakley, H. and Ferrie, J. (2016). Shocking Behavior: Random Wealth in Antebellum Georgia and Human Capital Across Generations *. *The Quarterly Journal of Economics*, 131(3):1455–1495.

Boneva, T. and Rauh, C. (2018). Parental beliefs about returns to educational investments—the later the better? *Journal of the European Economic Association*, 16(6):1669–1711.

Bonhomme, S., Holzheu, K., Lamadon, T., Manresa, E., Mogstad, M., and Setzler, B. (2020). How much should we trust estimates of firm effects and worker sorting? Technical Report 27368, National Bureau of Economic Research.

Borghans, L., Duckworth, A. L., Heckman, J. J., and Ter Weel, B. (2008). The economics and psychology of personality traits. *Journal of human Resources*, 43(4):972–1059.

Boserup, S. H., Kopczuk, W., and Kreiner, C. T. (2016). The role of bequests in shaping wealth inequality: evidence from danish wealth records. *American Economic Review*, 106(5):656–61.

Bound, J. and Solon, G. (1999). Double trouble: on the value of twins-based estimation of the return to schooling. *Economics of Education Review*, 18(2):169–182.

Bowles, S., Gintis, H., and Osborne, M. (2001). The determinants of earnings: A behavioral approach. *Journal of economic literature*, 39(4):1137–1176.

Branigan, A. R., McCallum, K. J., and Freese, J. (2013). Variation in the heritability of educational attainment: An international meta-analysis. *Social forces*, 92(1):109–140.

Bratberg, E., Davis, J., Mazumder, B., Nybom, M., Schnitzlein, D. D., and Vaage, K. (2017). A comparison of intergenerational mobility curves in Germany, Norway, Sweden, and the US. *The Scandinavian Journal of Economics*, 119(1):72–101.

Bratsberg, B., Røed, K., Raaum, O., Naylor, R., Ja¨ntti, M., Eriksson, T., and O¨sterbacka, E. (2007). Nonlinearities in intergenerational earnings mobility: consequences for cross-country comparisons. *The Economic Journal*, 117(519):C72–C92.

Bredtmann, J. and Smith, N. (2018). Inequalities in educational outcomes: How important is the family? *Oxford Bulletin of Economics and Statistics*, 80(6):1117–1144.

Carneiro, P., Garcia, I. L., Salvanes, K. G., and Tominey, E. (2021). Intergenerational mobility and the timing of parental income. *Journal of Political Economy*, 129(3):757–788.

Carneiro, P., Løken, K. V., and Salvanes, K. G. (2015). A flying start? Maternity leave benefits and long-run outcomes of children. *Journal of Political Economy*, 123(2):365–412.

Cascio, E. U. (2021). Early childhood education in the United States: What, when, where, who, how, and why. Technical report, National Bureau of Economic Research.

Cesarini, D., Lindqvist, E., Östling, R., and Wallace, B. (2016). Wealth, health, and child development: Evidence from administrative data on Swedish lottery players. *The Quarterly Journal of Economics*, 131(2):687–738.

Cesarini, D. and Visscher, P. M. (2017). Genetics and educational attainment. *npj Science of Learning*, 2(1):1–7.

Chetty, R. (2014). Improving opportunities for economic mobility in the United States.

Chetty, R., Friedman, J., Hendren, N., Jones, M., and Porter, S. (2018). The opportunity atlas: Mapping the childhood roots of social mobility. Technical Report 25147, National Bureau of Economic Research.

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.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014a). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Chetty, R., Hendren, N., Kline, P., Saez, E., and Turner, N. (2014b). Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *American Economic Review*, 104(5):141–47.

Collins, W. A., Maccoby, E. E., Steinberg, L., Hetherington, E. M., and Bornstein, M. H. (2000). Contemporary research on parenting: The case for nature and nurture. *American psychologist*, 55(2):218.

Conley, T. G. and Topa, G. (2002). Socio-economic distance and spatial patterns in unemployment. *Journal of Applied Econometrics*, 17(4):303–327.

Connor, D. S. and Storper, M. (2020). The changing geography of social mobility in the United States. *Proceedings of the National Academy of Sciences*, 117(48):30309–30317.

Cools, S., Fiva, J. H., and Kirkebøen, L. J. (2015). Causal effects of paternity leave on children and parents. *The Scandinavian Journal of Economics*, 117(3):801–828.

Cooper, K. and Stewart, K. (2017). Does money affect children's outcomes? An update. Technical Report 203, IZA Institute for Labor Economics.

Corak, M. (2006). *Do poor children become poor adults? Lessons from a cross-country comparison of generational earnings mobility*. Emerald Group Publishing Limited.

Corak, M. (2020). The canadian geography of intergenerational income mobility. *The Economic Journal*, 130(631):2134–2174.

Corak, M. and Piraino, P. (2011). The intergenerational transmission of employers. *Journal of Labor Economics*, 29(1):37–68.

Cornelissen, T., Dustmann, C., Raute, A., and Schönberg, U. (2018). Who benefits from universal child care? Estimating marginal returns to early child care attendance. *Journal of Political Economy*, 126(6):2356–2409.

Coulton, C., Crampton, D., Irwin, M., Spilsbury, J., and Korbin, J. (2007). How neighborhoods influence child maltreatment: A review of the literature and alternative pathways. *Child abuse & neglect*, 31:1117–42.

Cunha, F. and Heckman, J. (2007). The technology of skill formation. *American Economic Review*, 97(2):31–47.

Cunha, F., Heckman, J. J., Lochner, L., and Masterov, D. V. (2006). Interpreting the evidence on life cycle skill formation. *Handbook of the Economics of Education*, 1:697–812.

Cunha, F., Heckman, J. J., and Schennach, S. M. (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931.

Currie, J. and Almond, D. (2011). Human capital development before age five. In *Handbook of labor economics*, volume 4, pages 1315–1486. Elsevier.

Dahl, G. B., Kostøl, A. R., and Mogstad, M. (2014). Family welfare cultures. *The Quarterly Journal of Economics*, 129(4):1711–1752.

Dahl, G. B. and Lochner, L. (2012). The impact of family income on child achievement: Evidence from the earned income tax credit. *American Economic Review*, 102(5):1927–56.

Dahl, M. W. and DeLeire, T. (2008). *The association between children's earnings and fathers' lifetime earnings: estimates using administrative data*. University of Wisconsin-Madison, Institute for Research on Poverty.

De Haan, M. (2011). The effect of parents' schooling on child's schooling: a nonparametric bounds analysis. *Journal of Labor Economics*, 29(4):859–892.

Del Boca, D., Flinn, C., and Wiswall, M. (2014). Household choices and child development. *Review of Economic Studies*, 81(1):137–185.

Deutscher, N. and Mazumder, B. (2020). Intergenerational mobility across australia and the stability of regional estimates. *Labour Economics*, 66:101861.

Dhillon, A., Iversen, V., and Torsvik, G. (2021). Employee referral, social proximity and worker discipline: theory and suggestive evidence from india. *Cultural Change and Economic Development*.

Doepke, M. and Zilibotti, F. (2017). Parenting with style: Altruism and paternalism in intergenerational preference transmission. *Econometrica*, 85(5):1331–1371.

Drange, N. and Havnes, T. (2019). Early childcare and cognitive development: Evidence from an assignment lottery. *Journal of Labor Economics*, 37(2):581–620.

Durlauf, S. N. (2004). Chapter 50 neighborhood effects. In *Handbook of Regional and Urban Economics*, volume 4, pages 2173–2242. Elsevier.

Durlauf, S. N. and Ioannides, Y. M. (2010). Social interactions. *Annu. Rev. Econ.*, 2(1):451–478.

Dustmann, C. and Schönberg, U. (2012). Expansions in maternity leave coverage and children's long-term outcomes. *American Economic Journal: Applied Economics*, 4(3):190–224.

Elliott, J. R. (1999). Social isolation and labor market insulation:. *The Sociological Quarterly*, 40(2):199–216.

Engzell, P. and Tropf, F. C. (2019). Heritability of education rises with intergenerational mobility. *Proceedings of the National Academy of Sciences*, 116(51):25386–25388.

Eriksen, J. and Munk, M. D. (2020). The geography of intergenerational mobility – Danish evidence. *Economics Letters*, 189. Publisher: Elsevier.

Fagereng, A., Mogstad, M., and Rønning, M. (2021). Why do wealthy parents have wealthy children? *Journal of Political Economy*, 129(3):703–756.

Fauth, R. C., Leventhal, T., and Brooks-Gunn, J. (2004). Short-term effects of moving from public housing in poor to middle-class neighborhoods on low-income, minority adults' outcomes. *Social Science & Medicine*, 59(11):2271–2284.

Felfe, C. and Lalive, R. (2018). Does early child care affect children's development? *Journal of Public Economics*, 159:33–53.

Fort, M., Ichino, A., and Zanella, G. (2020). Cognitive and noncognitive costs of day care at age 0–2 for children in advantaged families. *Journal of Political Economy*, 128(1):158–205.

Galster, G. C. (2012). The mechanism(s) of neighbourhood effects : theory, evidence, and policy implications. In *Neighbourhood Effects Research: New Perspectives*, pages 23–56. Springer.

Galton, F. (1886). Regression towards mediocrity in hereditary stature. *The Journal of the Anthropological Institute of Great Britain and Ireland*, 15:246–263.

Gandil, M. (2019). Intergenerational mobility or inequality: What does rank-mobility measure? Unpublished. Available at https://sites.google.com/view/mikkelgandil/research?authuser=0.

Goldberger, A. S. (1979). Heritability. *Economica*, 46(184):327–347.

Hai, R. and Heckman, J. J. (2017). Inequality in human capital and endogenous credit constraints. *Review of economic dynamics*, 25:4–36.

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

Harding, D. J., Sanbonmatsu, L., Duncan, G. J., Gennetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R., Sciandra, M., and Ludwig, J. (2021). Evaluating contradictory experimental and non-experimental estimates of neighborhood effects on economic outcomes for adults. Technical Report 28454, National Bureau of Economic Research.

Havnes, T. and Mogstad, M. (2011a). Money for nothing? Universal child care and maternal employment. *Journal of Public Economics*, 95(11-12):1455–1465.

Havnes, T. and Mogstad, M. (2011b). No child left behind: Subsidized child care and children's long-run outcomes. *American Economic Journal: Economic Policy*, 3(2):97–129.

Havnes, T. and Mogstad, M. (2015). Is universal child care leveling the playing field? *Journal of public economics*, 127:100–114.

Heckman, J. J. (2005). The scientific model of causality. *Sociological methodology*, 35(1):1–97.

Heckman, J. J. (2007). The economics, technology, and neuroscience of human capability formation. *Proceedings of the national Academy of Sciences*, 104(33):13250–13255.

Heckman, J. J. and Karapakula, G. (2019). Intergenerational and intragenerational externalities of the perry preschool project. Technical Report 25889, National Bureau of Economic Research.

Heckman, J. J. and Landersø, R. (2020). Lessons from Denmark about inequality and social mobility. In *EALE SOLE AASLE World Conference*.

Heckman, J. J., Moon, S. H., Pinto, R., Savelyev, P. A., and Yavitz, A. (2010). The rate of return to the highscope perry preschool program. *Journal of public Economics*, 94(1-2):114–128.

Heckman, J. J. and Mosso, S. (2014). The economics of human development and social mobility. *Annu. Rev. Econ.*, 6(1):689–733.

Heidrich, S. (2017). Intergenerational mobility in Sweden: a regional perspective. *J Popul Econ*, 30(4):1241–1280.

Helsø, A.-L. (2021). Intergenerational income mobility in Denmark and the United States. *The Scandinavian Journal of Economics*, 123(2):508–531.

Herbst, C. M. (2013). The impact of non-parental child care on child development: Evidence from the summer participation "dip". *Journal of Public Economics*, 105:86–105.

Herbst, C. M. (2017). Universal child care, maternal employment, and children's long-run outcomes: Evidence from the US lanham act of 1940. *Journal of Labor Economics*, 35(2):519–564.

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.

Hyytinen, A., Ilmakunnas, P., Johansson, E., and Toivanen, O. (2019). Heritability of lifetime earnings. *The Journal of Economic Inequality*, 17(3):319–335.

Inoue, A. and Solon, G. (2010). Two-sample instrumental variables estimators. *The Review of Economics and Statistics*, 92(3):557–561.

Jantti, M., Bratsberg, B., Roed, K., Raaum, O., Naylor, R., Osterbacka, E., Bjorklund, A., and Eriksson, T. (2006). American exceptionalism in a new light: a comparison of inter-generational earnings mobility in the nordic countries, the United Kingdom and the United States. Technical Report 1938, IZA Institute for Labor Economics.

Jencks, C. and Mayer, S. E. (1990). Chapter 4 the social consequences of growing up in a poor neighborhood. In *Inner-City Poverty in the United States*, pages 111–186. National Academy Press.

Kalmijn, M. (1998). Intermarriage and homogamy: Causes, patterns, trends. *Annual review of sociology*, 24(1):395–421.

Karlson, K. and Landersø, R. (2021). The making and unmaking of opportunity: Educational mobility in 20th century-Denmark. Technical Report 14135, IZA Institute for Labor Economics.

Kirkeboen, L. J., Leuven, E., and Mogstad, M. (2016). Field of study, earnings, and self-selection*. *The Quarterly Journal of Economics*, 131(3):1057–1111.

Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2007.00733.x.

Knight, F. (1935). *The Ethics of Competition*. New York: Allan & Unwin.

Kong, A., Thorleifsson, G., Frigge, M. L., Vilhjalmsson, B. J., Young, A. I., Thorgeirsson, T. E., Benonisdottir, S., Oddsson, A., Halldorsson, B. V., Masson, G., et al. (2018). The nature of nurture: Effects of parental genotypes. *Science*, 359(6374):424–428.

Kramarz, F. and Skans, O. N. (2014). When strong ties are strong: Networks and youth labour market entry. *Review of Economic Studies*, 81(3):1164–1200.

Krueger, A. B. (2012). The rise and consequences of inequality in the United States. *Speech at the Center for American Progress*, 12.

Landersø, R. and Heckman, J. J. (2017). The Scandinavian fantasy: The sources of inter-generational mobility in Denmark and the US. *The Scandinavian journal of economics*, 119(1):178–230.

Lee, J. J., Wedow, R., Okbay, A., Kong, E., Maghzian, O., Zacher, M., Nguyen-Viet, T. A., Bowers, P., Sidorenko, J., Linnér, R. K., et al. (2018). Gene discovery and polygenic prediction from a genome-wide association study of educational attainment in 1.1 million individuals. *Nature genetics*, 50(8):1112–1121.

Lee, S. Y. and Seshadri, A. (2019). On the intergenerational transmission of economic status. *Journal of Political Economy*, 127(2):855–921.

Lentz, B. F. and Laband, D. N. (1989). Why so many children of doctors become doctors: Nepotism vs. human capital transfers. *Journal of Human Resources*, pages 396–413.

Lindbeck, A., Nyberg, S., and Weibull, J. W. (1999). Social norms and economic incentives in the welfare state. *The Quarterly Journal of Economics*, 114(1):1–35.

Liu, Q. and Skans, O. N. (2010). The duration of paid parental leave and children's scholastic performance. *The BE Journal of Economic Analysis & Policy*, 10(1).

Lochner, L. and Monge-Naranjo, A. (2016). Student loans and repayment: Theory, evidence, and policy. In *Handbook of the Economics of Education*, volume 5, pages 397–478. Elsevier.

Løken, K. V., Mogstad, M., and Wiswall, M. (2012). What linear estimators miss: The effects of family income on child outcomes. *American Economic Journal: Applied Economics*, 4(2):1–35.

Ludwig, J. (2012). Guest editor's introduction. *Cityscape*, 14(2):1–28. Publisher: US Department of Housing and Urban Development.

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):226–231.

Ludwig, J., Duncan, G. J., and Hirschfield, P. (2001). Urban poverty and juvenile crime: Evidence from a randomized housing-mobility experiment. *The Quarterly Journal of Economics*, 116(2):655–679.

Manski, C. F. (2000). Economic analysis of social interactions. *Journal of economic perspectives*, 14(3):115–136.

Manski, C. F. (2011). Genes, eyeglasses, and social policy. *Journal of Economic Perspectives*, 25(4):83–94.

Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data. *Review of Economics and Statistics*, 87(2):235–255.

Mazumder, B. (2008). Sibling similarities and economic inequality in the US. *Journal of Population Economics*, 21(3):685–701.

Mazumder, B. (2016). Estimating the intergenerational elasticity and rank association in the United States: Overcoming the current limitations of tax data. In *Inequality: Causes and consequences*. Emerald Group Publishing Limited.

Miller, P., Mulvey, C., and Martin, N. (2001). Genetic and environmental contributions to educational attainment in australia. *Economics of Education review*, 20(3):211–224.

Mitnik, P., Bryant, V., Weber, M., and Grusky, D. B. (2015). New estimates of intergenerational mobility using administrative data. *Statistics of Income working paper. Washington: Internal Revenue Service*.

Mitnik, P. A. and Grusky, D. B. (2020). The intergenerational elasticity of what? The case for redefining the workhorse measure of economic mobility. *Sociological Methodology*, 50(1):47–95.

Modalsli, J. (2017). Intergenerational mobility in Norway, 1865–2011. *The Scandinavian Journal of Economics*, 119(1):34–71.

Mogstad, M. (2017). The human capital approach to intergenerational mobility. *Journal of Political Economy*, 125(6):1862–1868.

Mogstad, M., Romano, J. P., Shaikh, A., and Wilhelm, D. (2020). Inference for ranks with applications to mobility across neighborhoods and academic achievement across countries. Technical Report 26883, National Bureau of Economic Research.

Narayan, A., Van der Weide, R., Cojocaru, A., Lakner, C., Redaelli, S., Mahler, D. G., Rama-subbaiah, R. G. N., and Thewissen, S. (2018). *Fair Progress?: Economic Mobility Across Generations Around the World*. World Bank Publications.

Nicolaou, N., Shane, S., Cherkas, L., Hunkin, J., and Spector, T. D. (2008). Is the tendency to engage in entrepreneurship genetic? *Management Science*, 54(1):167–179.

Nicoletti, C., Salvanes, K. G., and Tominey, E. (2020). Mothers working during preschool years and child skills. does income compensate? Technical Report 13079, IZA Institute for Labor Economics.

OECD (2018). *A Broken Social Elevator? How to Promote Social Mobility*. OECD Publishing.

Okbay, A., Beauchamp, J. P., Fontana, M. A., Lee, J. J., Pers, T. H., Rietveld, C. A., Turley, P., Chen, G.-B., Emilsson, V., Meddens, S. F. W., et al. (2016). Genome-wide association study identifies 74 loci associated with educational attainment. *Nature*, 533(7604):539–542.

Papageorge, N. W. and Thom, K. (2020). Genes, education, and labor market outcomes: evidence from the health and retirement study. *Journal of the European Economic Association*, 18(3):1351–1399.

Pekkarinen, T., Salvanes, K. G., and Sarvimäki, M. (2017). The evolution of social mobility: Norway during the twentieth century. *The Scandinavian Journal of Economics*, 119(1):5–33.

Pekkarinen, T., Uusitalo, R., and Kerr, S. (2009). School tracking and intergenerational income mobility: Evidence from the Finnish comprehensive school reform. *Journal of Public Economics*, 93(7-8):965–973.

Pinto, R. (2019). Noncompliance as a rational choice: A framework that exploits compromises in social experiments to identify causal effects. Working paper.

Plomin, R., DeFries, J. C., and Loehlin, J. C. (1977). Genotype-environment interaction and correlation in the analysis of human behavior. *Psychological bulletin*, 84(2):309.

Polderman, T. J., Benyamin, B., De Leeuw, C. A., Sullivan, P. F., Van Bochoven, A., Visscher, P. M., and Posthuma, D. (2015). Meta-analysis of the heritability of human traits based on fifty years of twin studies. *Nature genetics*, 47(7):702–709.

Raaum, O., Salvanes, K., and SÃžrensen, E. (2006). The neighbourhood is not what it used to be. *Economic Journal - ECON J*, 116:200–222.

Ridley, M. (2003). Nature via nurture: Genes, experience, and what makes us human. *New York*.

Røed, K. and Raaum, O. (2003). Administrative registers–unexplored reservoirs of scientific knowledge? *The economic journal*, 113(488):F258–F281.

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.

Sacerdote, B. (2011). Nature and nurture effects on children's outcomes: What have we learned from studies of twins and adoptees? In *Handbook of social economics*, volume 1, pages 1–30. Elsevier.

Schmitz, L. L. and Conley, D. (2017). The effect of Vietnam-era conscription and genetic potential for educational attainment on schooling outcomes. *Economics of education review*, 61:85–97.

Schnitzlein, D. D. (2014). How important is the family? Evidence from sibling correlations in permanent earnings in the USA, Germany, and Denmark. *Journal of Population Economics*, 27(1):69–89.

Shang, Q. (2014). Endogenous neighborhood effects on welfare participation. *Empirical Economics*, 47(2):639–667.

Sharkey, P. (2010). The acute effect of local homicides on children's cognitive performance. *PNAS*, 107(26):11733–11738. Publisher: National Academy of Sciences Section: Social Sciences.

Solis, A. (2017). Credit access and college enrollment. *Journal of Political Economy*, 125(2):562–622.

Solon, G. (1992). Intergenerational income mobility in the United States. *The American Economic Review*, pages 393–408.

Solon, G. (1999). Intergenerational mobility in the labor market. In *Handbook of labor economics*, volume 3, pages 1761–1800. Elsevier.

Solon, G. (2004). A model of intergenerational mobility variation over time and place. In Corak, M., editor, *Generational income mobility in North America and Europe*, chapter Chap. 2 in Generational Income Mobility in North America and Europe, pages 38–47. Cambridge, UK: Cambridge University Press.

Solon, G. (2018). What do we know so far about multigenerational mobility? *The Economic Journal*, 128(612):F340–F352.

Solon, G., Corcoran, M., Gordon, R., and Laren, D. (1991). A longitudinal analysis of sibling correlations in economic status. *Journal of Human resources*, pages 509–534.

Solon, G., Page, M. E., and Duncan, G. J. (2000). Correlations between neighboring children in their subsequent educational attainment. *Review of Economics and Statistics*, 82(3):383–392.

Stigler, G. J. and Sherwin, R. A. (1985). The extent of the market. *The Journal of Law and Economics*, 28(3):555–585. Publisher: The University of Chicago Press.

Stuhler, J. (2018). A review of intergenerational mobility and its drivers. Technical Report JRC112247, Joint Research Centre.

Taubman, P. (1976). The determinants of earnings. genetics, family, and other environments: A study of white male twins. *The American Economic Review*, 66(5):858–870.

Trzaskowski, M., Harlaar, N., Arden, R., Krapohl, E., Rimfeld, K., McMillan, A., Dale, P. S., and Plomin, R. (2014). Genetic influence on family socioeconomic status and children's intelligence. *Intelligence*, 42:83–88.

Vartanian, T. (1997). Neighborhood effects on AFDC exits: Examining the social isolation, relative deprivation, and epidemic theories. *Social Service Review - SOC SERV REV*, 71:548–573.

Zimmerman, D. J. (1992). Regression toward mediocrity in economic stature. *The American Economic Review*, pages 409–429.