THE INTERGENERATIONAL TRANSMISSION OF FAMILY-INCOME ADVANTAGES IN THE UNITED STATES

Pablo A. Mitnik
Stanford Center on Poverty and Inequality
(pmitnik@stanford.edu)

Victoria Bryant
Statistics of Income Division, Internal Revenue Service
(victoria.l.bryant@irs.gov)

Michael Weber
Statistics of Income Division, Internal Revenue Service
(michael.e.weber@irs.gov)

July, 2018

The first author gratefully acknowledges research support from the Russell Sage Foundation and the Pew Charitable Trusts. The opinions expressed in this article are solely those of the authors and do not represent the opinions of the Internal Revenue Service or the Stanford Center on Poverty and Inequality.

Abstract

Estimates of economic persistence and mobility in the United States, as measured by the intergenerational elasticity (IGE), cover a very wide range. Nevertheless, careful analyses of the evidence suggested until recently that as much as half, and possibly more, of economic advantages are passed on from parents to children. This "dominant hypothesis" was seriously challenged by the first-ever study of family-income mobility based on tax data (Chetty et al. 2014), which provided estimates of family-income IGEs indicating that only one third of economic advantages are transmitted across generations and claimed that previous highly influential IGE estimates based on administrative data were upward biased. Using a different tax-based dataset, this article provides estimates of family-income IGEs that strongly support the dominant hypothesis. The article also carries out a one-to-one comparison between IGEs estimated with the two tax-based datasets and shows that Chetty et al.'s estimates were driven downward by a combination of attenuation, lifecycle, selection and functional-form biases. Lastly, the article determines the exact relationship between parental-income inequality, economic persistence and inequality of opportunity, which leads to the twofold conclusion that in the United States at least half of income inequality among parents is transformed into inequality of opportunity for their children, and that there is a very high level of inequality of opportunity in the country compared to most other highly-developed countries.

Introduction

To what extent are economic advantages passed on from parents to children in the United States? What share of economic inequality among families persists from one generation to the next? How much economic mobility across generations there is? How far is the country from achieving the normative ideal of equality of opportunity in the economic realm? The intergenerational elasticity has been, by a large margin, the measure most often employed to answer these crucial questions.[1] The IGE of men's earnings has been extensively estimated, in most cases with survey data (see reviews by Solon 1999; Corak 2006; Mitnik et al. 2018) but also with information from the Social Security Administration (Mazumder 2005; Dahl and DeLaire 2008). The IGE of family income has also been estimated in the United States, but less often than that of earnings and, until very recently, exclusively with survey data.[2]

The available IGE estimates cover a very wide range. Nevertheless, careful analyses of the evidence accumulated over decades of research, which were very strongly influenced by the administrative-data results reported by Mazumder (2005), led over time to the *tentative* conclusion that U.S. income and (men's) earnings' IGEs are not smaller than 0.5. For instance, in his most recent appraisal of the literature, Solon (2008:4) contended that once all downward biases in the estimation of the IGE are considered "it becomes plausible that the intergenerational elasticity in the United States may well be as large as 0.5 or 0.6." Similarly, Black and Devereaux (2011:1495-1496) recently wrote in the *Handbook of Labor Economics* that "a reasonable guess is an IGE in the US of about 0.5 to 0.6." An IGE in this range entails that at least half of economic advantages are passed on from one generation to the next, which has often been interpreted as *suggesting* a very high level of inequality of opportunity in the country, both in absolute terms and compared to most other highly-developed countries (for cross-country comparisons of IGEs see, e.g., Corak 2013; Jäntti et al. 2006). We will refer to the view that at least half of economic advantages persist across generations, and that this is a very high level of persistence, as the "dominant hypothesis." Importantly, as the United States is also characterized by high income inequality, that descriptive hypothesis is consistent with the causal hypothesis that more inequality leads to lower mobility; the latter hypothesis has long been discussed by sociologists and economists, and in recent times has often been considered in relationship to the "Great Gatsby curve" showing a negative correlation between the Gini coefficient and the IGE across countries (e.g., Bloome 2015; Corak 2013; Jerrim and Macmilliam 2015; Mitnik et al. 2016; Solon 2004; Torche 2005; Western and Bloome 2011).

The publication, in a top economics journal, of the first-ever study of family-income mobility in the United States based on tax data (Chetty et al. 2014) has cast serious doubts on the dominant hypothesis. In their highly influential article, Chetty et al. (2014) argued that Mazumder's (2005) approach for dealing with missing parental data had led to upward-biased IGE estimates.[3] In addition, although Chetty et al. (2014) found out that their IGE estimates were nonrobust to the treatment of children who did not file taxes as adults, they nevertheless reported a preferred estimate of the family-income IGE that is as low as 0.34 (for men and women pooled).[4] This estimate indicates that about one third, rather than at least half, of economic advantages are passed on from parents to children. It also indicates much less economic persistence than what has been generally assumed to be the case since the publication of Mazumder (2005); while economic persistence in the United States has been deemed, for some time now, to be highest or close to highest among highly-developed countries, an IGE of about 0.34 would mean that this persistence is in fact very close to the average persistence across those countries.[5] As the survey data employed in mobility research are affected by a long list of problems and limitations that reduce the confidence we can place on the resulting IGE estimates (Mitnik et al. 2018:10-11; Schoeni and Wiemers 2015), Chetty et al.'s (2014) lower-end estimate of the income IGE with high-quality tax data, together with their criticism of Mazumder's (2005) administrative-data results, have seriously undermined the epistemic status of the dominant hypothesis.[6]

IGEs are defined in terms of long-run income variables. However, as it has been nearly always the case in the literature, Chetty et al. (2014) used short-run proxy variables to estimate them. Mobility scholars have long emphasized that many things may go wrong in this context (e.g., Harder and Solon 2006; Mazumder 2005; Mitnik and Grusky 2017), so it is very important to determine whether Chetty et al.'s IGE estimates are on the mark. If the share of economic advantages transmitted across generations were really close to one third, that would mean that the dominant hypothesis is simply mistaken; in turn, the tenability of the notion that there is a causal link between inequality and mobility would be significantly reduced.

In this article we contend that nothing of the sort is the case. Relying on several samples from the Statistics of Income Mobility (SOI-M) Panel (Mitnik et al. 2015), a different tax-based dataset than that used by Chetty et al. (2014), we provide IGE estimates that strongly support the view that at least half of economic advantages are transmitted from parents to children and show that Chetty et al.'s estimates were driven downward by a combination of attenuation, lifecycle, selection, and functional-form biases. By supplementing Chetty et al.'s (2014) estimates with additional IGE estimates based on aggregate tax-based statistics they have made publicly available, we are able to carry out a one-to-one comparison between IGEs estimated with our and Chetty et al.'s (2014) data. This comparison shows that estimates based on the two tax-based datasets (and associated methodological decisions) are systematically and markedly different and imply quite contrasting assessments of the extent to which inequalities among families are transmitted across generations. Moreover, the comparison is consistent with our contention that the estimates reported, or based on the data employed, by Chetty et al. (2014) are affected by the aforementioned biases. Strong evidence that this is the case is then furnished by (a) using the SOI-M Panel to generate a dataset replicating the key "bias-generating features" of Chetty et al.'s sample and methodological decisions, and showing that the estimates move in the expected direction in all cases and are quite close to Chetty et al.'s after all biases are introduced, and (b) computing "Shapley decompositions" (Shorrocks 2013) to quantify the specific contribution of each bias to the differences in results.

In addition, in this article we provide the first formal account of the relationship between cross-sectional inequality, economic persistence and inequality of opportunity.[7] This allows us to advance a novel and compelling justification for the interpretation of an IGE as the share of economic advantages or inequality transmitted across generations (or "share interpretation" of the IGE). It also allows us to make the relationship between the transmission of economic advantages and inequality of opportunity transparent, and to show that (a) at least half of income inequality among parents is transformed into inequality of opportunity for their children, and (b) given what we know about cross-sectional income inequality, high economic persistence in the United States straightforwardly *entails* (as opposed to just suggesting) a very high level of inequality of opportunity in the country compared to most other highly-developed countries.

Our article is most closely related to Mazumder (2016) and Mitnik et al. (2018). Like us, Mazumder (2016) criticizes Chetty et al.'s (2014) IGE estimates. However, he focuses exclusively on one of the four IGEs we examine (the constant IGE conventionally estimated in the literature), only considers two of the four biases we discuss (the lifecycle and attenuation biases), and relies on survey data affected by the problems and limitations mentioned earlier and very different, both in nature and in terms of the period they cover, from the tax data used by Chetty et al. (2014). In contrast, we use a data set unaffected by those problems and limitations and built from data sources very similar to those used by Chetty et al. (2014), measure family income in a very similar way to theirs, and focus on a time period adjacent to the one they consider (2010 compared to 2011-2012).[8] Our article and Mitnik et al. (2018) are complementary. Both provide tax-based IGE estimates consistent with the dominant hypothesis and use for this purpose the same sample. However, we provide here family-income IGE estimates for men and women pooled rather than by gender (as Mitnik et al. [2018] do), which allows us to carry out a straightforward comparison with Chetty et al.'s (2014) estimates and with other estimates based on their data. More crucially, while Mitnik et al. (2018) focus exclusively on their preferred IGE concept, the IGE of expected income, here we pay equal attention to that IGE and to the IGE of the geometric mean of income; this makes direct comparisons with the large number of results reported in the literature, both for the United States and for other countries, possible (as we explain later, the latter IGE concept is what has been unwittingly estimated in the mobility literature). Equally important, while Mitnik et al. (2018) do not explain in any detail why their estimates and Chetty et al.'s (2014) estimates differ, accounting for why the latter and our estimates differ is one of our main goals. Lastly, Mitnik et al. (2018) do not offer any account of the relationship between cross-sectional inequality, IGEs and inequality of opportunity, rely on the (less appealing) standard justification for the share interpretation of IGEs, and do not provide a rationale for the interpretation of high economic persistence in the U.S. in terms of inequality of opportunity as we do here.

We lead off the article by introducing a generic notion of IGE, the four specific IGEs that are relevant for our arguments, and our account of the relationship between cross-sectional inequality, economic persistence and inequality of opportunity. This is followed by a description of estimators and potential estimation biases and the strategies used to address the latter. Next, we introduce the data we use and explain why they can be expected to lead to better estimates than those obtained by Chetty et al. (2014). After that we make our empirical case. The last section discusses our results and distills the article's main conclusions.

Conditional income distributions, intergenerational curves and economic persistence

Questions about the transmission of economic advantages from parents to children, the persistence of inequality across generations, economic mobility and inequality of opportunity may be expressed as questions about the distribution of children's income (as adults) conditional on their parents' income. If the conditional distribution of children's income does not vary across parental incomes, then no transmission of economic advantage or economic persistence exists and there is "perfect mobility" and full equality of opportunity.[9] If, on the contrary, the children's conditional distributions become "better" as parental income increases, then advantages are passed on in some degree, economic status is a persistent property, mobility is imperfect and there is inequality of opportunity.[10]

Therefore, a possible approach to answering questions about persistence and inequality of opportunity is to compare full conditional distributions across levels of parental income and assess how those distributions change as parental income increases (e.g. Lefranc et al. 2009). Data constraints, however, make it difficult to estimate full conditional distributions with precision. Moreover, it is rather unwieldy to compare them even if they can be estimated. One way to deal with these two problems is to (a) summarize the information contained in full conditional distributions by using a measure of central tendency, for instance the arithmetic mean or expectation, (b) specify and estimate an "intergenerational curve" relating the selected measure of central tendency of children's income to parental income, and (c) summarize the information in that curve in a way that is relevant for the questions at hand.

Analyses based on IGEs carry out the last task by focusing on the slope of the intergenerational curve, with the curve defined in log-log space rather than in the space spanned by the income variables. The focus on the slope may be motivated as follows: A fully flat intergenerational curve indicates no economic persistence or inequality of opportunity (at least in terms of the selected measure of central tendency of income), while an increasing curve indicates the opposite, so mobility scholars found natural to think of steeper curves as indicating more economic persistence and less mobility.[11] But why to define the curve in log-log space? The reason is that, in log-log space, the slope has attractive properties that it doesn't have otherwise: It is invariant to proportional economic growth (i.e., economic growth that leads to proportional increases in all children's incomes), to changes in measurement units, and to changes in price levels (see Mulligan 1997:25 for related comments). These are important properties, among other things because they make meaningful comparisons across countries and times possible.

So far we have referred to "the" slope of the intergenerational curve in log-log space, but the slope may vary across levels of parental income. There is good evidence, however, that at the population level the curve is monotonically increasing (that is, it always increases when parental income increases), at least with the measures of central tendency relevant here (see Chetty et al. 2014: Online Appendix Fig. 1). Therefore, the intuition that a steeper slope indicates less economic mobility and more economic persistence is still valid as long as we switch our focus to the expected slope across values of parental income—or, what is the same, to the *average slope* of the curve, where this average is a weighted average with weights equal to the density of each parental-income value.[12] In fact, this notion also covers the case in which the slope of the intergenerational curve is constant in log-log space—that is, when the elasticity is constant—as in this case the average slope is of course equal to that constant slope.

The foregoing suggests a characterization of a generic IGE as the *average slope of an intergenerational curve defined in log-log space*, with the specific measure of central tendency employed in the curve giving rise to a specific IGE concept. Further, actual IGE estimates also depend on the functional form posited for the relationship between children's and parental income in log-log space. It follows than an IGE is always an average point elasticity, across values of parental income, of a measure of central tendency of children's income with respect to parental income. Selecting a measure of central tendency and a functional form specifies a particular IGE.

Intergenerational elasticities, share interpretations and inequality of opportunity

The four specific IGEs that are relevant for the comparisons at the core of this article are obtained as indicated in Figure 1, where the measures of central tendency considered are the expectation and the geometric mean while the functional forms are a straight line and an unknown smooth function. As the assumption of a straight line naturally leads to the estimation of parametric constant-elasticity models, while the assumption of a smooth curve naturally leads to the estimation of nonparametric models, we will refer to the four IGEs as the constant IGE_e, the constant IGE_g, the nonparametric IGE_e, and the nonparametric IGE_g of children's family income with respect to parental income (where the subscripts "e" and "g" distinguish between IGE concepts, i.e., the IGE of the expectation and the IGE of the geometric mean).

Let's unpack these four IGEs, starting with the constant IGE_g. Its use of the geometric mean as the measure of central tendency is the unintended—and, until very recently, unnoticed—result of the reliance on logarithmically-transformed income variables to produce the elasticity estimates widely reported in the mobility literature (Mitnik and Grusky 2017). Indeed, the standard population regression function (PRF) posited by mobility scholars (e.g., Solon 1999) is:

$$E(\ln Y\,|x) = \beta_0 + \beta_1 \ln x, \qquad [1]$$

which may be written as

$$\ln GM\,(Y|x) = \beta_0 + \beta_1 \ln x, \qquad [1']$$

where Y is the children's long-run income, X is long-run parental income, GM is the geometric mean operator and β_1 is the income IGE the mobility literature has ubiquitously estimated (see Online Appendix B).[13] This conventionally estimated elasticity has been widely misinterpreted: While mobility scholars have assumed that they estimated the elasticity of the *expectation* of children's income, they in fact estimated, as Equation [1′] shows, the elasticity of the *geometric mean* of children's income —i.e., β_1 is the percentage differential in the geometric mean of children's long-run income with respect to a marginal percentage differential in parental long-run income.

Due to a host of conceptual and methodological problems, the constant IGE_g is not an attractive estimand; crucially, estimation of β_1 with the data that are in most cases available can be expected to be affected by a well-understood form of selection bias (Mitnik and Grusky 2017). As we will show later, this problem strikes with a vengeance if the data employed are tax data including a substantial share of nonfiler children and these (or a large share of them) are dropped from the analysis. Despite its shortcomings, the constant IGE_g is the workhorse measure of mobility employed in the literature. For this reason, it is important to include it in our comparison.

Estimation of the constant IGE_e is based on the following PRF:

$$\ln E(Y|x) = \alpha_0 + \alpha_1 \ln x, \qquad [2]$$

where α_1 is the percentage differential in the expectation of children's long-run income with respect to a marginal percentage differential in parental long-run income. Mitnik and Grusky (2017) have called for making this elasticity the workhorse intergenerational elasticity. They have shown that this elasticity is what mobility scholars assumed they were obtaining by estimating Equation [1], and that all interpretations incorrectly applied to β_1 are valid or approximately valid for α_1.

A key such interpretation, which we strongly emphasize in this article, is the share interpretation, according to which an IGE measures the share of economic advantages or inequality transmitted across generations. This interpretation may be based on two different analyses—both for the IGE_g and for the IGE_e—although only the analysis we present first has been previously considered in the literature (see, e.g., Mitnik and Grusky 2017). In the case of the IGE_e, from Equation [2] it immediately follows that:

$$\alpha_1 = \frac{\ln E(Y|x_2) - \ln E(Y|x_1)}{\ln x_2 - \ln x_1}, \qquad [3]$$

where we assume, without any loss of generality, that $x_2 > x_1$. As a difference in logarithms approximates well a percent difference as long as the latter is fairly small, it is the case that

$$\alpha_1 \approx \frac{E(Y|x_2) - E(Y|x_1)}{E(Y|x_1)} \left[\frac{x_2 - x_1}{x_1}\right]^{-1} \qquad [4]$$

as long as the ratio between x_2 and x_1 is not much larger than one. That is, under the conditions just specified, α_1 is approximately equal to the ratio between the proportional difference in the expected income of children and the proportional difference in the income of their parents. In order to interpret α_1 as the share of advantages or inequality that is transmitted across generations requires invoking a local notion of advantage or inequality, i.e., advantage or inequality between families that are close in the income distribution, and to measure this advantage or inequality by the proportional difference between their incomes.

The second, novel, analysis uses a regular (i.e., global) measure of income inequality and therefore maps more smoothly into the notion that income inequality is transmitted across generations and that IGEs measure the extent to which that is the case. Crucially, this analysis also makes the relationship between IGEs and inequality of opportunity fully transparent. Denoting the standard deviation operator by SD, it follows from Equation [2] that:

$$\alpha_1 = \frac{SD(\ln E(Y|X))}{SD(\ln X)}. \qquad [5]$$

Now (a) the SD of the logarithm of an income variable is a commonly-used measure of income inequality (e.g., Bourguignon and Morrisson 2002), and (b) indexing children's income opportunities by their conditional expected incomes is by far the most common approach in the bourgeoning empirical literature on inequality of opportunity (e.g., Ferreyra and Gignoux 2011; Brunori et al. 2013). It follows from Equation [5] that α_1 is equal to the ratio between the inequality in children's income opportunities, or just inequality of opportunity, and the inequality in parental income.[14]

Equations [4] and [5] both underlie the share interpretation of the constant IGE_e, according to which the latter measures the share of income inequality or advantages among families that is passed on to (the expected incomes of) their children. Analogous analyses can be provided for the constant IGE_g by simply replacing expectations by geometric means (and $\alpha_1$by β_1) in equations [3], [4] and [5], which leads to the conclusion that that elasticity measures the share of income inequality or advantages among families that is passed on to (the geometric mean of the incomes of) their children. Here, for the connection to inequality of opportunity to be maintained, children's opportunities need to be indexed not by their expected incomes conditional on parental income—as proposed in the literature on inequality of opportunity—but by the corresponding conditional geometric means.[15]

Equation [5] and its counterpart for the IGE_g make transparently clear that a constant IGE does not measure inequality of opportunity per se, as the latter is equal to the inequality among parents multiplied by the IGE; in other words, *an IGE measures the rate at which parental-income inequality gets transformed into inequality of opportunity*. However, because countries with larger IGEs also tend to exhibit more cross-sectional income inequality (as reflected in the popular Great Gatsby Curve), economic persistence and inequality of opportunity are highly (but far from perfectly) correlated (Brunori, et al. 2013).[16]

The assumption of a constant IGE has been adopted more as a matter of necessity (given the small samples available) than by virtue of any strong prior that it in fact holds. Unfortunately, if it doesn't hold, then estimates obtained under the constant-elasticity assumption are affected by *functional-form bias* (e.g., Bratberg et al. 2007; see also Corak and Heisz 1999). To address this potential bias, Mitnik et al. (2018) proposed estimating the nonparametric IGE_e in the bottom-left of Figure 1. Here, the assumption that the curve relating children's expected income to their parental income is a straight line in log-log space is replaced by the following, much weaker, assumption:

$$\ln \text{E}(Y|x) = \text{F}(\ln \text{x}), \qquad [6]$$

where F is an unknown smooth function. The resulting persistence measure has a share interpretation that generalizes the first one we discussed in the case of the constant IGE_e.

Assume that pairs of families whose incomes do not differ much in percent terms are randomly drawn from the parental-income distribution. Then, the nonparametric IGE_e approximates the expected share of inequality or advantages passed on to their children, across all possible such random draws. (If the constant-elasticity assumption holds, then this interpretation also applies, trivially, to the constant IGE_e, see Equation [4].) In addition, we show in Online Appendix C that the nonparametric IGE_e also provides an approximation to the ratio between the global measures of inequality of opportunity and of parental income we introduced above (the quantity in the right-hand side of Equation [5]).

The last IGE is the nonparametric IGE_g. Its definition and interpretation are analogous to those of the nonparametric IGE_e, but substituting the conditional geometric mean of children's income for their conditional expectation and, accordingly, replacing [6] by:

$$\ln GM(\mathrm{Y}|\mathrm{x}) = \mathrm{E}(\ln Y\,|x) = \mathrm{G}(\ln x), \qquad [7]$$

where G is an unknown smooth function. A share interpretation analogous to that advanced in the case of the nonparametric IGE_e is of course available.

Before finishing this section, it seems important to stress what the ontological status of IGEs is. Sometimes causal language slips in in discussions of IGE estimates, and there obviously are causal processes underlying them, but the IGEs themselves are not causal parameters. Rather, these are all descriptive measures—akin to, for instance, the Gini coefficient. Although they are all-important measures, it is simply a category mistake to interpret them as measures of the causal effects of parental income. This in turn entails that the characterization of the IGE_g and the IGE_e as, respectively, "person-weighted" and "dollar-weighted" elasticities (Chetty et al, 2014: 1574 and Online Appendix C), according to which the IGE_g is a simple average of personlevel *behavioral* elasticities while the IGE_e is a weighted average of the same elasticities that gives more weight to people with more income, is invalid (see Mitnik 2017c for details).

Estimators and potential biases due to the use of short-run income measures

As measures of long-run (e.g., lifetime) income are almost never available, estimation of IGEs is typically carried out by substituting short-run proxy variables for the long-run variables of interest. All empirical estimates we discuss in later sections were obtained by (a) replacing the long-run income of children (Y) by an annual family-income measure pertaining to when the children were in their 30s, or by an average of such annual measures over two years, and (b) replacing the long-run income of parents (X) by their average income over several years, pertaining to when the children were "young" (we provide details on the exact income measures used later). Below, when we refer to Equations [1], [2], [6] and [7], we are referring to versions of these equations with the short-run proxy variables substituted for their long-run counterparts. We first introduce the relevant estimators and then discuss the potential biases that the substitution of short-run proxy variables may generate.

Estimators

The estimators employed to estimate the four IGEs in Figure 1 are depicted in Figure 2. We briefly describe them here. A detailed discussion of these estimators as well as other estimation issues can be found in Online Appendix D.

Following the most common approach in the mobility literature, all estimates of the constant IGE_g we discuss are the result of estimating the PRF of Equation [1] by OLS, i.e., of employing the "OLS log-log estimator." In contrast, the estimates of the constant IGE_e rely on two different approaches. The new estimates we present here were obtained by estimating Equation [2] with the Poisson Pseudo Maximum Likelihood (PPML) estimator (Santos Silva and

Tenreyro 2006). The estimate of the constant IGE_e reported by Chetty et al. (2014) we will discuss is based on a two-step estimator of the same equation. In the first step, nonparametric estimates of $\ln E(Y | \ln x)$ are generated; in the second step, an estimate of α_1 is obtained by running an OLS regression of the estimates of $\ln E(Y | \ln x)$ on the corresponding $\ln x$ values.

The estimators of the nonparametric IGEs are two-step estimators in all cases: The first step produces nonparametric estimates of a number of points in the relevant intergenerational curve—i.e., the curve defined by either Equation [6] or Equation [7]—while the second step estimates the average slope of the curve through a numerical approximation based on the estimated points. Across datasets, the estimators only differ on the nonparametric approach used to estimate the points of the intergenerational curves and on the number of points that are estimated and employed in the numerical approximations.

Potential biases

When estimating the constant IGE_g, substituting short-run proxy measures for the long-run measures of interest opens the door to three biases, two of which have been extensively discussed in the mobility literature. First, measurement error produces substantial *attenuation bias* if annual measures of parental income, or other measures based on a few years of information, are used to estimate Equation [1] by OLS (e.g., Solon 1999; Mazumder 2005).[17] Second, as income-age profiles differ across economic origins, *lifecycle biases* result from using proxy measures taken when parents or children are too young or too old to represent lifetime differences well (e.g., Black and Devereux 2011). A formal joint analysis of these two biases is provided by the generalized error-in-variables (GEiV) model (Haider and Solon 2006). It follows from this model that using measures of economic status pertaining to specific ages should eliminate the bulk of the lifecycle biases—while the evidence available suggests that using parents' and children's information close to age 40 is the best approach (Haider and Solon 2006; Böhlmark and Lindquist 2006; Mazumder 2001; Mitnik 2017b; Nybon and Stuhler 2016). To address the problem of attenuation bias, the GEiV model—and many analyses predating it (e.g. Solon 1992)—suggests using parents' average income over several years as the measure of parental income. There is strong evidence that the bias can be substantially reduced this way if the average is computed over enough years, although there is disagreement on how many years are necessary to eliminate most of it (see Mitnik et al. 2018:9-15; see also our Online Appendix F).

Mitnik and Grusky (2017) showed that a third bias looming over the estimation of the constant IGE_g is *selection bias*. Mobility scholars have addressed what they have perceived as the practical problem of the logarithm of zero being undefined with the expedient of dropping children with zero income from samples.[18] As a result, estimation of that IGE with short-run proxy measures typically involves the use of a "censored sample," which generates a well-understood form of selection bias (e.g., Heckman 2008); as a large share of children have zero short-run income, the magnitude of this bias may be substantial (see Mitnik and Grusky 2017 for details).[19] Mitnik and Grusky (2017) have shown that there is no attractive work-around for this problem.[20]

The IGE_e is immune to the selection bias affecting the IGE_g, as its estimation with short-run proxy measures does not require dropping children with zero income from samples. At the same time, Mitnik (2017a) advanced and empirically validated a generalized error-in-variables model (the GEiVE model) indicating that the use of proxy measures makes estimation of the constant IGE_e with the PPML estimator vulnerable to lifecycle and attenuation biases very similar to those affecting estimation of the constant IGE_g with the OLS estimator. He also showed that the same strategies employed with the IGE_g to eliminate, or at least greatly reduce, those biases, can be expected to be effective when estimating the constant IGE_e.

Neither a formal measurement model, nor empirical evidence on the methodological issues at hand, are available for the estimation of the nonparametric IGEs with short-run proxy variables. Nevertheless, we expect that the same biases affecting estimation of the constant IGEs will be at play in the case of the nonparametric IGEs.

Data and variables

The SOI-M Panel, described in detail by Mitnik et al. (2015), is based on tax returns and other administrative data (e.g., W-2 and 1099 forms). It represents all children born between 1972 and 1975 who were living in the United States in 1987, and includes parental income information collected when the children were between 15 and 23 years old and children's income information, starting at age 26, for the period 1998-2010. Almost all empirical results we present here are based on the SOI-M Panel. In our analyses, we use information on children's income pertaining to 2010, when they were 35-38 years old, and to 2004, when they were 29-32 years old.

Our exclusive concern in this paper is with IGEs of (pre-tax) family income. While the income measures employed in our analyses are annual measures in the case of children (either for 2010 or for 2004), they are averages over several years in the case of parents. We use a measure of parental income based on nine years of parental information (pertaining to when the children were 15 to 23 years old) as well as a five-year measure (pertaining to when the children were 15 to 19 years old).

We exclude from our analyses children with (a) negative income, (b) income over $7,000,000, (c) more than two years (in the case of the five-year measure) or three years (in the case of the nine-year measure) of missing parental information, (d) nonpositive average parental income, or (e) average parental income over $7,000,000. Depending on which income measures (for both parents and children) are used, these sample selection rules generate four samples that differ slightly in size and demographic composition (children's mortality between 2004 and 2010 also plays a minor role). Using an obvious nomenclature, we refer to them as the 2004-5y, 2004-9y, 2010-5y and 2010-9y samples. Descriptive statistics for the four samples are shown in Tables 1 and 2. Table 1 reports the number of observations and the gender and age of the children included in each sample, the origin of their income information and the number of missing years of parental information among those retained. Table 2 shows the weighted mean and standard deviation of the income variables, for children and parents, and of parental age (the SOI-M Panel is based on a stratified random sample of 1987 tax returns, so all our analyses employ sampling weights). The income variables are expressed in 2010 dollars using the Consumer Price Index for Urban Consumers - Research Series (CPI-U-RS).

For some more limited purposes, we also use aggregate statistics that Chetty et al. (2014) have made publicly available (see the sources of Table 3). The microdata underlying those statistics—which are the microdata used by Chetty et al. (2014) in their research—represent the birth cohorts 1980-1982. Here, the children's income is their average income in 2011-2012, when they were between 29 and 32 years old, while their parental income is measured by averaging five years of information, when they were between 14 and 20 years old. Income refers to (pre-tax) family income in both cases.

Key differences in samples and methodological decisions and their expected effects on estimates

There are four differences between the samples used and the methodological decisions made in Chetty et al.'s (2014) research and in ours that can be expected to generate differences in estimates. The first three pertain to the ages of the children relied on to produce preferred estimates, the number of years of parental information employed to construct parental-income measures for the same purpose, and the treatment of children without tax or other administrative information on their income. The fourth difference, which is of a different nature, involves decisions on the summary mobility measures that are important to estimate given overall research goals. We discuss these differences and their expected effects on estimates in turn.

Children's ages and parental income measures

A central goal of the research by Chetty et al. (2014) was to study mobility within quite small geographic areas. To this end they relied on the full population of tax records—which are only available starting in 1996—to construct the very large dataset (i.e., a dataset with close to ten million observations) employed in their core analyses. In contrast, here our focus is on the correct estimation of national-level IGEs, and to this end we rely on the SOI-M 2010-9y sample, which has close to 12,500 observations. For reasons that will become clear later, we refer to this sample as the "SOI-M best sample."

Figure 3 allows to compare this sample to the data employed by Chetty et al (2014). The figure makes apparent two differences that are of central interest in the light of our previous discussion of lifecycle and attenuation biases. First, while the children in Chetty et al.'s data are in their early 30s when their income is measured, those in the SOI-M best sample are in their late 30s, i.e., much closer to the age the literature has deemed optimal. This suggests that Chetty et al.'s estimates may be substantially downward biased while ours should be much less affected by lifecycle bias. Importantly, although Chetty et al. (2014) emphatically denied that their IGE estimates were significantly impinged by this bias, in agreement with Mazumder (2016) we find the evidence they provided to support their claim flawed (see Online Appendix E).

Second, Chetty et al.'s measure of parental income is based on five years of information while the corresponding measure in the SOI-M best sample is based on nine years. In his very influential article, Mazumder (2005) argued that up to 16 years of information are needed to eliminate or nearly eliminate attenuation bias; although this is likely to be an overestimate (see Online Appendix F), there is a rather broad consensus that five years of information are not enough to secure good estimates of the (constant) IGE_g. Similarly, Mitnik (2017a) reported that approximately 13 years of information are needed to eliminate the bulk of attenuation bias when estimating the (constant) IGE_e with survey data (although, for reasons discussed in Online Appendix F, it is very likely that fewer years are needed with administrative data). This suggests that Chetty et al.'s IGE estimates may be significantly reduced by attenuation bias, while ours should be much less affected. Although Chetty et al. (2014) strongly rejected that their estimates of the constant IGE_g are affected by attenuation bias, their evidence for their claim that five years of parental information are enough to eliminate the bulk of that bias is quite weak (see Online Appendix G).

"Nonadmin" children

The third key difference concerns the treatment of nonfiler children without other administrative income information. In any tax year a number of people do not file taxes, mostly because their incomes are below the thresholds that make filing mandatory. Not surprisingly, then, not all children included in Chetty et al.'s (2014) data and in the SOI-M best sample filed taxes in the tax years when their income was measured (2011-2012 and 2010, respectively).[21] To address this problem, both Chetty et al. (2014) and Mitnik et al. (2015)—the latter, when building the SOI-M Panel—resorted to other administrative sources (e.g., earnings from W-2 forms, unemployment-insurance income from 1099 forms) to approximate the income of some nonfiler children. For other nonfiler children, however, alternative administrative information was not available. As a result, 6.1 percent of children in Chetty et al.'s core sample (2014: Online Appendix Table III) and 7.1 percent of children in the SOI-M best sample (Table 1) are nonfilers without other administrative information. In what follows, we refer to this subset of nonfiler children as "nonadmin children."

Although nonadmin children may have some income—mostly from work in the informal economy and from transfers from sources not covered in constructing the income measures, e.g., Temporary Assistance to Needy Families (TANF)—Chetty et al. (2014) opted for assigning them an income of zero. This can be expected to generate an *upward* bias in the estimation of the IGE_e, as it underestimates the income of nonadmin children, whose share decreases sharply as parental income increases (see, e.g., Chetty et al. 2014: Figure 1). The notable robustness of the IGE_e to the treatment of nonadmins (Mitnik et al. 2018:30-33) indicates, however, that this bias should be small. In contrast, the estimation of the IGE_g becomes very problematic under this approach. The reason is that it requires dropping nonadmin children, which in turn means that the resulting estimates can be expected to be seriously affected by the selection bias discussed by Mitnik and Grusky (2017). Chetty et al. did acknowledge that dropping nonadmin children from the sample "overstates the degree of intergenerational mobility" (Chetty et al. 2014:1573). However, it is clear that they deemed any ensuing selection bias small—in particular, small enough to make comparisons between the resulting estimates and previous estimates in the literature perfectly meaningful. Indeed, not only did Chetty et al. (2014) make estimates obtained by dropping nonadmin children they preferred estimates, but they also argued that they were "broadly consistent with previous results, with the exception of Mazumder's (2005) and Clark's

(2014) IGE estimates, which imply much lower levels of intergenerational mobility" (Chetty et al. 2014:1558).[22]

Instead of assigning an income of zero to nonadmin children, we resort to data from the Annual Social and Economic Supplement of the Current Population Survey (CPS-ASEC) to carry out mean imputation. The CPS-ASEC data include information on "likely nonfilers," who are identified using a tax simulation model developed by the U.S. Census Bureau. Using this information, we compute the mean income of likely nonadmin children separately for each of six gender-age groups in the SOI-M Panel (see Online Appendix H). Then, under the assumption that the expectation of nonadmin children's income is independent of parental income, we assign those mean values to the corresponding nonadmin children and use the resulting income variables to estimate the IGE_e.

Implementing the mean-imputation strategy with the IGE_g is less straightforward. The reason is that approximately one-third of CPS likely nonfilers without earnings or UI income—or "CPS nonadmins," for short—have zero family income; as in this context it's necessary to impute the mean of the logarithm of income (rather than mean income), CPS nonadmins with zero reported income pose a problem. The mobility literature suggests two approaches for addressing it. The first approach is to assume that zero reported income is the result of a mechanism unrelated to true income, and therefore that those with zero income may be unproblematically dropped when computing average log income. This is, of course, equivalent to the assumption mobility scholars have almost always made, implicitly, when estimating the IGE_g (and which generates the downward selection bias discussed by Mitnik and Grusky [2017]). The second approach is to assume that zero reported income does not reflect true income but is nevertheless indicative of very low positive income; this is the assumption implicitly made in those very few cases in which the standard approach for estimating the IGE_g was deemed problematic, and which typically manifested itself in the assignment of an income of one dollar to children with zero reported income (e.g., Couch and Lillard 1998). Using these two approaches to compute the mean log income of CPS nonadmins within each gender-age group, we generate a set of upper and a set of lower imputation values, respectively, and employ them to conduct mean imputation (see Online Appendix H). As we have independent evidence (for the constant IGE_g) that the estimates that result bracket the true value of the elasticity, we interpret them as lower- and upper-bound estimates, respectively, of the IGE_g.[23]

Mean imputation is a methodologically superior approach than zero imputation. Although it may be expected to make a small difference for IGE_e estimates—which, as already indicated, are very robust to the treatment of nonadmin children—it can be expected to make a substantial difference for the estimation of the IGE_g.

Summary mobility measures

Chetty et al. (2014) and Mitnik et al. (2018) have argued that intergenerational curves are markedly nonlinear in log-log space. Moreover, the evidence they offered—especially the nonparametric curves Chetty et al. (2014) estimated, which are based on binned scatter plots relying on millions of observations—is rather conclusive. In contrast, the constant IGE conventionally estimated by mobility scholars assumes a linear relationship. As it should be apparent, however, a constant IGE is a poor summary measure of economic mobility and persistence when the true curve is far from linear (Bratsberg et al. 2007).

As Chetty et al.'s (2014) main goal wasn't the estimation of the share of economic inequality transmitted across generations in the United States but the study of geographic variability in mobility within the country, their response to the finding that the IGE_g is markedly nonlinear (and to the estimation difficulties generated by nonadmin children) was to switch to a different measure not affected by those problems for the bulk of their analyses.[24] This strategy was justified in the light of their main goal. Here, however, our main goal requires the estimation of national-level IGEs while the finding of marked nonlinearities requires that we estimate nonparametric IGEs. With respect to our goal, estimates based on the constant-elasticity assumption are affected by functional-form bias. Therefore, in the specific sense just discussed, the IGE estimates reported by Chetty et al. (2014) should be affected by functional-form bias. The presence of this bias should contribute to explain the differences between those estimates and our preferred IGE estimates.

Empirical analyses

We have claimed that Chetty et al.'s (2014) IGE estimates should be affected by lifecycle, attenuation and functional-form biases and, more generally, that IGE estimates based on their data should be affected by the first two biases. We have also claimed that IGE_g estimates based on their data should be affected by selection bias while IGE_e estimates should be affected by what we may refer as "imputation bias" (due to the imputation of zero income to nonadmin children). Lastly, we have argued that IGE estimates based on the SOI-M sample should be essentially unaffected by selection and imputation biases, and not much affected by lifecycle and attenuation biases; it follows that those estimates should be substantially higher than estimates based on Chetty et al.'s data.[25] We now make the empirical case for our claims. We start by providing baseline estimates of both the IGE_e and the IGE_g of family income, based on our and on Chetty et al.'s (2014) data. Next, we resort to a sample from the SOI-M Panel that allows us to replicate the key bias-generating features of Chetty et al.'s data and show that it leads to IGE estimates than are much lower than those obtained with the SOI-M best sample and quite close to those based on Chetty et al.'s data. After that we conduct Shapley decompositions (Shorrocks 2013) to quantify the specific contribution of each bias. All estimates pertain to men and women pooled.

Baseline estimates

The top panel of Table 3 presents baseline IGE_e figures. The first column shows Chetty et al.'s (2014) estimate of the constant IGE_e and our estimate of the nonparametric IGE_e using Chetty et al.'s data. These estimates are 0.34 and 0.38, respectively. The second column shows our estimates of the IGE_e based on the SOI-M best sample. At 0.46 and 0.50, these estimates are 35 percent (constant IGE_e) and 32 percent (nonparametric IGE_e) larger than the corresponding estimates based on Chetty et al.'s data. Moreover, our nonparametric estimate—that is, our preferred estimate—based on the SOI-M best sample is almost 50 percent larger than the IGE_e estimate reported by Chetty et al. (2014) under the constant-elasticity assumption (0.50 compared to 0.34).

The bottom panel of Table 3 presents baseline IGE_g figures. Chetty et al.'s (2014) estimate of the constant IGE_g, as well as our estimate of the nonparametric IGE_g using Chetty et al.'s data, are again in the first column. As just noted, these estimates were obtained by dropping nonadmin children. The next two columns show our lower- and upper-bound estimates based on the SOI-M best sample. These estimates put the constant IGE_g in the 0.46-0.67 range and the nonparametric IGE_g in the 0.53-0.74 range. The corresponding estimates based on Chetty et al.'s (2014) data are 0.34 and 0.39. This means that even our lower-bound estimates are about 35 percent larger than the corresponding estimates based on Chetty et al.'s data and methodological decisions. Moreover, our lower-bound nonparametric estimate is a full 56 percent larger than the IGE_g estimate reported by Chetty et al. (2014) under the constant-elasticity assumption (0.53

compared to 0.34). The difference is, of course, much larger when we consider our upper-bound estimate.

Baseline estimates compared to estimates based on "replication data" from the SOI-M Panel

The results of Table 3 are consistent with our hypothesis that Chetty et al.'s (2014) data lead to downward-biased IGE_e estimates. To provide evidence for this hypothesis, we next resort to a SOI-M sample replicating the key bias-generating features of Chetty et al.'s data. If we are on the mark, we should produce with this sample much lower IGE_e estimates than those obtained with the SOI-M best sample.

Figure 4 shows how we approximate two features of Chetty et al.'s data that, we have argued, generate (or exacerbate) lifecycle and attenuation biases. First, we use the children's family-income information for 2004, when they were 29-32 years old (the same ages of the children in Chetty et al.'s sample). Second, like Chetty et al., we employ a parental-income measure based on five years of information. The resulting sample is the 2004-5y sample we described earlier. In addition, to further mirror Chetty et al.'s approach, we substitute zeros for the CPS-based mean income values imputed to nonadmin children in the SOI-M data. We refer to the resulting sample as the "SOI-M all-biases sample."

Figure 5 shows the IGE_e estimates produced with this sample, together with the corresponding estimates from Table 3. At 0.38 and 0.41, the new estimates of the constant and nonparametric IGE_e are substantially smaller than those obtained with the best sample and—particularly for the latter estimate—quite close to those based on Chetty et al.'s data. The results in Figure 5 suggest that between two thirds (constant IGE_e) and three quarters (nonparametric IGE_e) of the within-model differences in estimates are accounted by lifecycle, attenuation and imputation biases affecting the estimates based on Chetty et al.'s data.[26] Likewise, the results suggest that, depending on the computation method used, the aforementioned biases plus functional-form bias account for between three quarters and better than four fifths of the difference between (a) our nonparametric (i.e., preferred) IGE_e estimate based on the SOI-M best sample, and (b) the IGE_e estimate reported by Chetty et al. (2014) under the constant-elasticity assumption.[27] The residual differences that are left unexplained are most likely due to sampling variability, the discrepancy in periods (2004 versus 2011-2012), the discrepancy in cohorts (1972-1975 versus 1980-1982), and small differences in the way in which variables are computed, sample inclusion rules, and other methodological decisions.

One such difference—which, we show below, cannot be ignored in the case of the IGE_g—is that Chetty et al. (2014) employed a measure of children's income based on two years of information. This, however, does not contribute to accounting for the residual difference in IGE_e estimates. Using a sample similar to the SOI-M all-biases sample but in which the children's income measure is their average income in 2003-2004 (which, for future reference, we call the "SOI-M two-year all-biases sample"), our estimates are unchanged at the second-decimal level we are reporting them.

To provide evidence for our hypothesis regarding the differences in IGE_g estimates, we also estimate this IGE with data replicating the key bias-generating features of Chetty et al.'s (2014) data and methodological decisions. As in the case of the IGE_e, if our hypotheses are correct, these estimates should be much lower than those obtained with the SOI-M best sample. Here, unlike in the case of the IGE_e, whether we use the SOI-M all-biases sample or its two-year counterpart should make a difference. The reason is that we expect fewer children with zero income in a two-year than in a one-year sample (as a child needs to be nonadmin in both years for his or her income to be zero in the former sample but only in one year in the latter sample).

As children with zero income are dropped from the analysis when we mirror Chetty et al.'s (2014) approach, the SOI-M all-biases sample should overstate selection bias and therefore total bias. For future reference, we present estimates of the IGE_g using both the two-year and the one-year SOI-M all-biases samples.

Figure 6 shows the results, together with the corresponding estimates based on Chetty et al.'s (2014) data and the corresponding lower- and upper-bound estimates based on the SOI-M best sample. The estimates of the constant and nonparametric IGE_g generated with the SOI-M two-year all-biases sample are 0.30 and 0.40, respectively, which means that they are much smaller than (even) the lower-bound estimates obtained with the SOI-M best sample. They are also quite close to the estimates based on Chetty et al.'s data—in fact the nonparametric estimates are almost identical. And, as anticipated, they are slightly larger (on average, 7.6 percent larger) than the estimates based on the SOI all biases sample, which relies on only one year of children's information and therefore overstates selection bias.

Figure 6 suggests that between two thirds (constant IGE_g) and better than nine tenths (nonparametric IGE_g) of the within-model lower-bound differences in estimates are accounted by lifecycle, attenuation and selection biases affecting the estimates based on Chetty et al.'s data.[28] Similarly, the figure suggests that, depending on the computation method used, these three biases plus functional-form bias account for between fourth fifths and nineteen twentieths of the difference between (a) the (constant) IGE_g estimate reported by Chetty et al. (2014), and (b) our lower-bound estimate of the nonparametric IGE_g.[29] As before, the residual differences that are left unexplained are most likely due to sampling variability, the discrepancy in periods and cohorts, and small differences in methodological decisions. Of course, those residual differences become much smaller in proportional terms (if not essentially nil), if instead of comparing Chetty et al.'s (2014) estimates with our lower-bound estimates we compare them with our upper-bound estimates. For instance, the four biases account for between 90 and 97.5 percent of the difference between the constant IGE_g estimate reported by Chetty et al. (2014) and our upper-bound nonparametric IGE_g estimate.[30]

Shapley decompositions

In Figure 5 we introduced, for each model, all bias-generating features at once, which does not allow to compute the size of each individual bias, or even provide evidence that *all* biases we posited in the case of the IGE_e are really at play and have the expected signs. We use the 2004-5y, 2004-9y, 2010-5y and 2010-9y samples described earlier to identify the direction and magnitude of each individual bias. For this purpose, we take estimates based on the 2010-9y sample to be free of lifecycle and attenuation biases, those based on the 2004-5y sample to be affected by both biases, and those based on the other two samples to be each affected by one bias but not the other. In each of the four samples, we may either keep the CPS-based mean imputation for nonadmin children's income or replace it by \$0 imputation (so as to "capture" the bias generated by the latter imputation). And, in each case, we may estimate the IGE_e under the constant-elasticity assumption or estimate it nonparametrically (so as to capture functional-form bias). There are therefore 16 possible combinations of sample, imputation and model, each of which can be used to obtain an IGE_e estimate affected by a specific set of biases.[31] Comparing the estimates produced with any two of such combinations, differing only in that one is affected by a bias that the other avoids, allows us to compute the marginal effect of introducing the feature generating that bias (given a common set of other bias-generating features). The size of each bias, however, cannot be determined by simply introducing all bias-generating features in some arbitrary order, and computing the difference in estimates that results from adding each one in turn, as the effect of introducing each bias-generating feature depends on which others are already present. For this reason, we do instead a Shapley decomposition (Shorrocks 2013) of the total bias, whereby the size of each bias is defined as *its average marginal contribution to total bias across all possible orders in which the four bias-generating features can be introduced*. In such Shapley decomposition, the sum of the individual biases is guaranteed to be equal to the total bias.

Table 4 shows the results of computing model-specific Shapley decompositions of total bias. As expected, the lifecycle and attenuation biases are downward biases (so their contributions to total bias are positive), while the imputation bias is an upward bias (so its contribution to total bias is negative). For both the constant and nonparametric estimates, the contribution of the lifecycle bias to total bias is about twice as large as that of the attenuation bias. And although the imputation bias is much smaller in absolute value than the other two biases, it still is the case that the use of zero imputation rather than mean imputation reduces total bias by about 15 percent.[32]

Table 5 shows the results of the across-models Shapley decomposition of total bias, which is defined as the difference between the nonparametric estimate based on the SOI-M best sample and the constant-elasticity estimate based on the SOI-M all-biases sample. The signs of the contributions of the attenuation, lifecycle and imputation biases are the same as in the model-specific decompositions, while the contribution of the functional-form bias is positive, indicating that it is a negative bias (as expected). As in Table 4, the contribution of the lifecycle bias is the largest—about twice as large as those of the attenuation and functional-form biases. Here, the imputation bias reduces total bias by about eight percent.[33] If, for the sake of simplicity, we ignore this bias, then the lifecycle bias contributes about half of the total bias while the attenuation and functional-form biases contribute about one quarter each.

The foregoing results provide strong evidence that the bulk of the difference between Chetty et al.'s (2014) constant IGE_e estimate and our nonparametric IGE_e estimate based on the SOI-M best sample (see Table 3), is due to the combined effects of lifecycle, attenuation, imputation and functional-form biases affecting the former estimate, with lifecycle bias making the largest contribution to that difference. Importantly, in Online Appendix I we show that our key attenuation-bias result so far—i.e., that five years of parental information are not enough to eliminate the bulk of attenuation bias in the case of the IGE_e—is robust to alternative ways of examining this bias. We also provide evidence suggesting that our IGE_e estimates are (nearly) free of attenuation bias.

Computing Shapley decompositions of total bias is more complicated in the case of the IGE_g. Indeed, to properly capture the effects of lifecycle, attenuation, functional-form and selection bias on the estimates based on Chetty et al.'s (2014) data and methodological decisions, all estimates should be obtained with two-year samples. However, for reasons explained in Online Appendix J, proceeding this way is not feasible. As an alternative, we do the following. As we did in the case of the IGE_e, for each estimate of total bias we compute 16 estimates, one for each possible combination of sample, imputation and model, using in all cases one-year samples. But before computing the Shapley decompositions, we adjust upward all estimates based on samples where the CPS-based mean imputations for nonadmin children were replaced by zeros, so as to reflect the fact that one-year samples overestimate selection bias. After introducing these adjustments, we compute the Shapley decompositions as before. Importantly, with this approach the total biases are exactly what is wanted in all cases (see Online Appendix J for details, and the notes to Tables 6 and 7 for the definitions of the total biases of interest).

The top panel of Table 6 presents the model-specific Shapley decompositions of total bias with respect to our lower-bound estimates. As expected, the selection bias is a downward bias (so its contribution to total bias is positive). The lifecycle and attenuation biases also have their expected negative signs, as was the case with the IGE_e. Unlike with the IGE_e, however, the relative contributions of the attenuation and lifecycle biases vary significantly across models. The bottom panel of Table 6 presents the same decompositions but with respect to our upper-bound estimates, with all biases exhibiting the expected signs. Also as expected, both the absolute size and the relative importance of selection bias are now markedly larger in both models (selection bias is more than five times as large in size and more than twice as large in relative importance in both cases). The relative importance of attenuation bias is smaller and that of lifecycle bias is much smaller (i.e., disproportionally smaller), in both models.

The top panel of Table 7 shows the results of the across-models Shapley decomposition of total bias with respect to our lower-bound estimates. As in the case of the IGE_e, all biases—including functional-form bias—are downward biases. However, while in that case lifecycle-cycle bias accounted for roughly half of total bias and the attenuation and functional-form biases accounted for roughly one quarter each, in the case of the IGE_g functional-form bias accounts for roughly two-fifths of total bias while lifecycle, attenuation and selection bias account for roughly one fifth each. As expected, both the absolute and relative importance of selection bias increases markedly in the Shapley decomposition with respect to our upper-bound estimate, which is shown in the bottom panel of the same table. Here a full three-fifths of the difference is accounted for by selection bias, while the functional-form and attenuation biases account for most of the rest.

The foregoing results provide strong evidence that the bulk of the differences between Chetty et al.'s (2014) constant IGE_g estimates and our preferred (i.e., nonparametric) IGE_g estimates, which we presented in Table 3, are due to the combined effects of lifecycle, attenuation, selection and functional-form biases affecting the former estimates. Functional-form bias makes the largest contribution to those differences in the case of the lower-bound estimates, while selection bias makes the largest contribution in the case of the upper-bound estimates.

Like with the IGE_e, in Online Appendix I we show that our second key attenuation-bias result—i.e., that five years of information are not enough to eliminate the bulk of attenuation bias in IGE_g estimates—is robust to alternative ways of examining this bias. Unlike with the IGE_e, however, we also show that although using nine years of information to estimate the IGE_g reduces that bias substantially, it is not likely to (nearly) eliminate it. This provides further support for our claim that Chetty et al.'s (2014) IGE_g estimates greatly understate true economic persistence as measured by that IGE concept.

Discussion and conclusions

Assessing the extent to which economic advantages are transmitted across generations is a quite difficult endeavor. In the United States, IGE estimates cover a very wide range and are mostly based on survey data affected by many problems and limitations. Nevertheless, the balance of the evidence suggested until recently that as much as half, and possibly more, of economic advantages are passed on from parents to children. This view, which was greatly influenced by Mazumder's (2005) administrative-data estimates, was seriously challenged when Chetty et al. (2014) (a) advanced a plausible argument to the effect that Mazumder's estimates were upward biased, and (b) provided tax-data estimates indicating that only one third of family-income advantages are transmitted across generations. Moreover, by suggesting that economic persistence in the United States is very close to the average persistence across highly-developed countries, Chetty et al.'s estimates also undermined the prima-facie tenability of the notion that there is a causal link between inequality and mobility.

Using a different tax-based data set better suited for the estimation of national-level IGEs, and a better approach for dealing with the "nonadmin problem," we have provided new estimates of family-income IGEs (for men and women pooled). Our preferred nonparametric estimates put the IGE of expected income at 0.5 and the conventional IGE estimated in the literature in the 0.53-0.74 range (while the corresponding constant-elasticity estimates are about ten percent lower).

We have also shown that Chetty et al.'s (2014) estimates of family-income IGEs and, more generally, IGE estimates based on their data and methodological decisions are downwardly biased due to the joint effects of attenuation, lifecycle, selection (in the case of the IGE_g) and functional-form biases. The evidence we have provided, which relies on tax-based samples mirroring the bias-generating features of Chetty et al.'s tax data and methodological decisions, is very strong. We used those samples both to generate estimates reproducing the bulk of the overall biases in the estimates based on Chetty et al.'s data and decisions and to isolate the absolute and relative contributions of each individual bias.

We first focused on the IGE of expected income. Our results indicated that (a) between two thirds and three quarters of the within-model differences between our estimates and estimates based on Chetty et al.'s (2014) data are accounted by lifecycle, attenuation and imputation biases affecting the latter, and (b) between three quarters and four fifths of the difference between our nonparametric estimate and the constant-elasticity estimate reported by Chetty et al. (2014) is accounted by the same three biases plus functional-form bias. In all cases, lifecycle bias is the main driver behind the observed differences, attenuation bias and (when relevant) functional-form bias play important roles, and imputation bias makes the differences somewhat smaller (as this bias is positive rather than negative).

We next focused on the IGE concept conventionally considered in the literature. Our analyses indicated that (a) between two thirds and better than nine tenths of the within-model differences between our lower-bound estimates and estimates based on Chetty et al.'s (2014) data are accounted by lifecycle, attenuation and selection biases, and (b) between fourth fifths and nineteen twentieths of the difference between our lower-bound nonparametric estimate and the constant-elasticity estimate reported by Chetty et al. (2014) is accounted by the same three biases plus functional-form bias. Here the main drivers of the observed differences vary across analyses, selection bias is substantial in all cases, and functional-form bias (when relevant) plays a larger role than in the case of the IGE of expected income. When we examined instead the differences with respect to our upper-bound estimates, the share of the differences accounted for by the four biases we have considered, as well as the relative importance of selection bias, increased substantially.

We have used high-quality tax data to show that, in the United States, at least half of economic advantages are transmitted across generations. In contrast with the estimates based on Chetty et al.'s (2014) data, the estimates based on the SOI-M best sample fully support the dominant hypothesis. In the light of the novel share interpretation advanced in this paper, our estimates indicate that in the United States at least half of income inequality among parents is transformed into inequality of opportunity for their children, regardless of whether we use the conditional expectation or the conditional geometric mean to define intergenerational curves and to index children's income opportunities. Our results also indicate that that share—that is, the rate at which income inequality is transformed into inequality of opportunity—may well be substantially larger than half when the conditional geometric mean is chosen as the opportunity index (as the mobility literature has implicitly done, at least until very recently).

Our results have further implications for inequality of opportunity in a comparative perspective. We have demonstrated that inequality of opportunity is equal to the product between parental-income inequality and economic persistence. Family-income inequality in the United States is largest among highly-developed countries (e.g., Gornick and Milanovic 2015), so it's safe to assume that the country also ranks close to the top, if not at the very top, in terms of parental-income inequality. Given that our empirical results offer strong support to the notion that economic persistence in the United States is close to highest, if not highest, among such countries (i.e., the dominant hypothesis), it immediately follows that there is a very high level of inequality of opportunity in the country compared to most other highly-developed countries.

Notes

[1] The IGE has almost always been estimated by assuming that it is constant across levels of parental income. Although strictly speaking a measure of the persistence of economic differences across generations (e.g., Jäntti et al. 2006:8), or economic persistence for short, the constant IGE has been commonly interpreted as a measure of economic mobility as well (in which a high IGE signifies low mobility).

[2] The IGE of women's earnings has been seldom estimated, in part because women's earnings are typically not considered to be a meaningful measure of their overall economic status (Chadwick and Solon 2002:335). One of the motivations for estimating the IGE of family income is that it allows to circumvent this limitation.

[3] See Online Appendix A for our discussion of Chetty et al.'s (2014) argument and Mazumder's (2016) response. Our main conclusions are that (a) Chetty et al. advanced a plausible empirical hypothesis, and (b) the evidence available does not allow to establish whether it is on the mark or not. See also Online Appendix F, where we conclude that Mazumder's (2005) estimates of attenuation bias in IGE estimation—as opposed to his IGE estimates—are most likely upward biased, as Chetty et al.'s (2014) also argued.

[4] Chetty et al. (2014) used a measure of positional mobility and persistence, the rank-rank slope, to conduct the bulk of their analyses. Here we are only concerned with their IGE estimates.

[5] After excluding the United States, the average IGE among the countries in Corak's (2013: Figure 1) Great Gatsby Curve, which only includes highly advanced economies, is 0.30.

[6] Chetty et al's 0.34 estimate has been included among the "best current estimates" of the IGE (Fox et al. 2016: 535) and some researchers have even interpreted it as a new benchmark for

today's United States (Lee and Seshadri forthcoming: 2 and 5; see also Yum 2015; Daruich and Kozlowski 2016).

[7] Or, more precisely, the first such an account that is correct; see note 15.

[8] Mazumder (2016) draws extensively on our own arguments here, which appeared in much less developed form in a working paper that is a predecessor to this article (and which circulated with a different title).

[9] Of course, the equality of opportunity alluded here is with respect to parental income only, taken *for the sake of the analysis* to be the only factor outside of people's control, that is, the only "circumstance" (see, e.g., Roemer 1998).

[10] The children's conditional distributions become better as parental income increases if, for instance, they are ordered by a relation of first-order stochastic dominance.

[11] Often, steeper curves are also interpreted as indicating more inequality of opportunity, which is only correct if the inequality in parental income is the same across curves. More on this later.

[12] From here on, whenever we refer to the "average slope" of the curve, it should be understood that we are referring to this weighted average.

[13] For Z and W any scalar random variables, we use expressions like "Z$|w$" as a shorthand for "$Z|W = w$."

[14] Mutatis mutandis, the same caveat of note 9 applies here. It has been shown that inequality of opportunity with respect to parental income (or any other specific circumstance or incomplete set of circumstances) provides a lower bound to total inequality of opportunity (Ferreyra and Gignoux 2011: Sec. 3).

[15] That is: $\beta_1 = \frac{SD(\ln GM(Y|X))}{SD(\ln X)}$ and $SD(\ln GM(Y|X))$ is a measure of inequality of opportunity as long as children's income opportunities are indexed by their conditional geometric means.

Lefranc et al. (2009) claimed to have provided a (different) formal account of the relationship between the conventionally estimated IGE (the constant IGE_g) and inequality of opportunity. That account is flawed because it wrongly assumes that the parameter β_0 in Equation [1] is a structural parameter allowing for a counterfactual interpretation.

[16] As many before them, Brunori, et al. (2013:13) claimed that, conceptually, economic persistence and inequality of opportunity "should be very closely related" but did not provide an account of that relationship.

[17] In this context, measurement error is defined with respect to the long-run measure. It therefore includes the effects of volatility over time in short-run income.

[18] The underlying problem is, of course, that the geometric mean (the unwittingly selected measure of central tendency) is undefined when a variable includes zero in its support.

[19] That there is a large share of children with zero short-run (e.g., annual) individual earnings is well known. The same is true for family income: 5.4 and 8.0 percent of adults ages 29-32 had zero family income in 2011-2012, according to data from the Current Population Survey and the American Community Survey, respectively (Chetty et al. 2014: Online Appendix Table IV).

[20] In particular, keeping all observations and replacing zeros by a "small amount" does not provide a solution, as estimates vary widely depending on the exact amount that is substituted.

[21] The problem is generated by children that do not file at all in the corresponding tax year, regardless of whether they file by the filing deadline or not. Both Chetty et al.'s (2014) and the SOI-M Panel data do include tax-based income information of late filers.

[22] The list of previous estimates of the constant IGE_g inconsistent with Chetty et al.'s (2014) estimates is substantially longer than they contend. Focusing only on family-income IGE estimates equal or larger than 0.5 and published before 2014, the list should include Abul Naga

(2001 and 2002), Bratsberg et al. (2007), Chadwick and Solon (2002), Jantti et al. (2006), and Hertz (2005 and 2006). Two post-2014 survey-data studies based on many years of parental information have reported IGE_g estimates substantially larger than 0.5 (Mazumder 2016; Mitnik 2017b).

[23] This evidence is based on left-censored parametric models that assume that the children's income follows the Dagum three-parameter distribution (e.g., Kleiber and Kotz 2003) and in which the income of nonadmins is taken to be below the tax filing thresholds. We will report these results in a separate paper. Relying on this type of models for our purposes here was not feasible given our central interest in comparing nonparametric estimates obtained using the SOI-M Panel and Chetty et al.'s (2014) data.

[24] See note 4.

[25] Strictly speaking, this only follows in the case of the IGE_g. In the case of the IGE_e, it follows if we further assume that the imputation bias is relatively small (as this bias is an upward bias). As we mentioned earlier, there is empirical evidence that this is the case.

[26] In the case of the constant IGE_e, the difference falls from 0.12 (0.46 – 0.34) with the best sample to 0.04 (0.38 – 0.34) with the all-biases sample, that is, only one-third of the original difference remains; in the case of the nonparametric IGE_e, the difference falls from 0.12 (0.50 – 0.38) to 0.03 (0.41 – 0.38), that is, only one quarter remains.

[27] We can proceed in two ways to compute the share of the observed difference accounted for by the four biases. The difference of 0.16 (0.50 – 0.34) (a) falls to 0.04 (0.38 – 0.34) once the SOI-M all-biases sample is employed to compute the constant IGE_e, or (b) falls to 0.03 (0.41 – 0.38) once the SOI-M all-biases sample is employed to compute the nonparametric IGE_e, and the estimate of the constant IGE_e based on Chetty et al.'s (2014) data is replaced by an estimate of

the nonparametric IGE_e based on the same data. The shares accounted for by the four biases are respectively, three quarters (1 – 0.04/0.16) or more than four fifths (1 - 0.03/0.16).

[28] In the case of the constant IGE_g, the difference in estimates falls from 0.12 (0.46 – 0.34) with the best sample to 0.04 (0.34 – 0.30) with the two-year all-biases sample, that is, only one-third of the original difference remains; in the case of the nonparametric IGE_g, the difference falls from 0.14 (0.53 – 0.39) to 0.01 (0.40 – 0.39), that is, only seven percent remains.

[29] We can proceed in two ways to compute the share of the observed difference accounted for by the four biases. The difference of 0.19 (0.53 – 0.34) (a) falls to 0.04 (0.34 – 0.30) once the SOI-M two-year all-biases sample is employed to compute the constant IGE_g, or (b) falls to 0.01 (0.40 – 0.39) once the SOI-M two-year all-biases sample is employed to compute the lower-bound of the nonparametric IGE_g, and the estimate of the constant IGE_g based on Chetty et al.'s (2014) data is replaced by an estimate of the nonparametric IGE_g based on the same data. The shares accounted for by the four biases are, respectively, almost four-fifths (1 – 0.04/0.19) or almost 95 percent (1 - 0.01/0.19).

[30] The difference of 0.40 (0.74 – 0.34) falls to 0.04 (0.34 – 030) or to 0.01 (0.40 – 0.39), depending on the computation method (see previous note). The shares accounted for by the four biases are, respectively, 90 percent (1 – 0.04/0.40) or 97.5 percent (1 - 0.01/0.40).

[31] This set is empty in the case of the nonparametric estimate based on the 2010-9y sample with CPS-based mean imputation for nonadmin children (the SOI-M best sample). It only includes all four biases we have considered in the case of the constant-elasticity estimate based on the 2004-5y sample with $0 imputation for nonadmin children (the SOI-M all-biases sample).

[32] In the case of the constant IGE_e, 0.08 / (0.08 + 0.01) = 0.85; in the case of the nonparametric IGE_e, 0.09 / (0.09 + 0.01) = 0.86.

[33] That is, 0.12 / (0.12 + 0.01) = 0.92.

References

Abul Naga. Ramses. 2001. "Galtonian Regression of Intergenerational Income Linkages: Biased Procedures, a New Estimator and Mean-Square Error Comparisons." STICERD - Distributional Analysis Research Programme Papers 53, London School of Economics.

Abul Naga, Ramses. 2002. "Estimating the Intergenerational Correlation of Incomes: An Errors-in-Variables Framework." *Economica* 69: 69-91.

Black, Sandra and Paul Devereux. 2011. "Recent Developments in Intergenerational Mobility." *Handbook of Labor Economics*, Volume 4b, edited by David Card and Orley Ashenfelter. Amsterdam: Elsevier.

Bloome, Deirdre. 2015. "Income Inequality and Intergenerational Income Mobility in the United States." *Social Forces*, 93(3): 1047-1080.

Böhlmark, Anders and Matthew Lindquist. 2006. "Life-Cycle Variations in the Association between Current and Lifetime Income: Replication and Extension for Sweden," *Journal of Labor Economics*, 24(4): 879–896.

Bourguignon, François and Christian Morrisson. 2002. "Inequality among World Citizens: 1820-1992." *The American Economic Review* 92 (4): 727-744.

Bratsberg, Bernt, Knut Røed, Oddbjørn Raaum, Robin Naylor, Markus Jäntti, Tor Eriksson and Eva Österbacka. 2007. "Nonlinearities in Intergenerational Earnings Mobility. Consequences for Cross-Country Comparisons." *The Economic Journal*, 117 (March) C72-C92.

Brunori, Paolo, Francisco Ferreira, and Vito Peragine. 2013. "Inequality of Opportunity, Income Inequality and Economic Mobility: Some International Comparisons." Policy Research Working Paper 6304, World Bank, Development Research Group.

Chadwick, Laura and Gary Solon. 2002. "Intergenerational Income Mobility among Daughters." *The American Economic Review* 92 (1): 335-344.

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

Clark 2014. *The Son Also Rises: Surnames and the History of Social Mobility*. Princeton, NJ: Princeton University Press.

Corak, Miles. 2006. "Do Poor Children Become Poor Adults? Lessons from a Cross Country Comparison of Generational Earnings Mobility." IZA Discussion Paper No. 1993.

Corak, Miles. 2013. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility." *Journal of Economic Perspectives* 27(3): 79-102.

Corak, Miles and Andrew Heisz. 1999. "The Intergenerational Income Mobility of Canadian Men." *Journal of Human Resources* 34(3): 504-33.

Couch, Kenneth, and Dean Lillard. 1998. "Sample Selection Rules and the Intergenerational Correlation of Earnings." *Labour Economics* 5: 313-329.

Dahl, Molly and Thomas DeLeire. 2008. “The Association between Children’s Earnings and Fathers’ Lifetime Earnings: Estimates Using Administrative Data.” Institute for Research on Poverty Discussion Paper 1342-08, University of Wisconsin-Madison.

Daruich, Diego and Julain Kozlowski. 2016. "Explaining Income Inequality and Intergenerational Mobility: The Role of Fertility and Family Transfers." Working Papers 2018-11, Federal Reserve Bank of St. Louis.

Ferreyra, Francisco and Jérémie Gignoux. 2011. “The Measurement of Inequality of Opportunity: Theory and an Application to Latin America.” *Review of Income and Wealth* 57 (4): 622-657.

Fox, Liana, Florencia Torche, and Jane Waldfogel. 2016. “Intergenerational Mobility.” In *Oxford Handbook of Poverty and Society*, edited by David Brady and Linda Burton. Oxford: Oxford University Press.

Gornick, Janet and Branko Milanovic. 2015. “Income Inequality in the United States in Cross-National Perspective: Redistribution Revisited.” Luxemburg Income Study Center Research Brief 1/2015.

Haider, Steven and Gary Solon. 2006. “Life-Cycle Variation in the Association between Current and Lifetime Earnings.” *American Economic Review* 96(4): 1308-1320.

Hertz, Tom. 2005. “Rags, Riches and Race: The Intergenerational Economic Mobility of Black and White Families in the United States.” In *Unequal Chances. Family Background and Economic Success*, edited by Samuel Bowles, Herbert Gintis, and Melissa Osborne Groves. New York, Princeton and Oxford: Russell Sage and Princeton University Press.

Hertz, Tom. 2006. “Understanding Mobility in America.” Report by the Center for American Progress.

Jäntti, Markus, Bernt Bratsberg, Knut Røed, Oddbjørn Raaum, Robin Naylor, Eva Österbacka, Anders Björklund, and Tor Eriksson. 2006. “American Exceptionalism in a New Light: A Comparison of Intergenerational Earnings Mobility in the Nordic Countries, the United Kingdom and the United States.” IZA Discussion Paper No. 1938.

Jerrim, John and Lindsey Macmilliam. 2015. “Income Inequality, Intergenerational Mobility, and the Great Gatsby Curve: Is Education the Key?” *Social Forces* 94(2): 505–533.

Kleiber, Christian and Samuel Kotz. 2003. *Statistical Size Distributions in Economics and Actuarial Sciences*. Hoboken, NJ: Wiley.

Lee, Sang Yoon and Ananth Seshadri. Forthcoming.” On the Intergenerational Transmission of Economic Status.” *Journal of Political Economy*. DOI: 10.1086/700765.

Lefranc, Arnaud, Nicolas Pistolesi and Alain Trannoy. 2009. “The Link between Inequality of Opportunity for Income Acquisition and Income Inequality: The French Example 1977-1993.” Available at https://sites.google.com/site/lefrancarnaud/Home/research.

Mazumder, Bhashkar. 2001. "The Miss-measurement of Permanent Earnings: New Evidence from Social Security Earnings Data." Federal Reserve Bank of Chicago Working Paper 2001-24.

Mazumder, Bhashkar. 2005. "Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data." *The Review of Economics and Statistics* 87(2): 235-255.

Mazumder, Bhashkar. 2016. "Estimating the Intergenerational Elasticity and Rank Association in the United States: Overcoming the Current Limitation of Tax Data." In *Inequality: Causes and Consequences*, edited by Lorenzo Cappellari, Solomon Polacheck, and Konstantinos Tatsiramos. Bingley: Emerald.

Mitnik, Pablo. 2017a. "Estimating the Intergenerational Elasticity of Expected Income with Short-Run Income Measures: A Generalized Error-in-Variables Model." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo. 2017b. "Intergenerational Income Elasticities, Instrumental Variable Estimation, and Bracketing Strategies." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo. 2017c. "'Person-Weighted' versus 'Dollar-Weighted': A Flawed Characterization of Two Intergenerational Income Elasticities." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo, Victoria Bryant, Michael Weber and David Grusky. 2015. "New Estimates of Intergenerational Mobility Using Administrative Data." SOI Working Paper, Statistics of Income Division, Internal Revenue Service.

Mitnik, Pablo, Victoria Bryant, Michael Weber and David Grusky. 2018. "A Very Uneven Playing Field: Economic Mobility in the United States." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo A., Erin Cumberworth, and David B. Grusky. 2016. "Social Mobility in a High-Inequality Regime." *Annals of the American Academy of Political and Social Science*, 663(1): 140-184.

Mitnik, Pablo and David Grusky. 2017. The Intergenerational Elasticity of What? The Case for Redefining the Workhorse Measure of Economic Mobility." Stanford Center on Poverty and Inequality Working Paper.

Mulligan, Casey. 1997. *Parental Priorities and Economic Inequality*. Chicago: University of Chicago Press.

Torche, Florencia. 2005. "Unequal but Fluid: Social Mobility in Chile in Comparative Perspective." *American Sociological Review* 70(3):422-450.

Nybom, Martin and Jan Stuhler. 2016. "Heterogeneous Income Profiles and Life-Cycle Bias in Intergenerational Mobility Estimation." *The Journal of Human Resources* 15(1): 239-268.

Roemer 1998. *Equality of Opportunity*. Cambridge, Mass.: Harvard University Press.

Santos Silva, J. M. C. and Silvna Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 88(4): 641-658.

Schoeni, Robert and Emily Wiemers. 2015. “The Implications of Selective Attrition for Estimates of Intergenerational Elasticity of Family income.” *Journal of Economic Inequality* 13(3): 351–372.

Shorrocks, Anthony. 2013. “Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value.” *The Journal of Economic Inequality* 11(1): 99-26.

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

Solon, Gary. 1999. “Intergenerational Mobility in the Labor Market.” *Handbook of Labor Economics*, Volume 3A, edited by Orley C. Ashenfelter and David Card. Amsterdam: Elsevier.

Solon, Gary. 2004. “A Model of Intergenerational Mobility Variation over Time and Place.” In *Generational Income Mobility in North America and Europa*, edited by Miles Corak. Cambridge: Cambridge University Press.

Solon, Gary. 2008. “Intergenerational Income Mobility.” *The New Palgrave Dictionary of Economics,* Second Edition, edited by Steven Durlauf and Lawrence Blume. Basingstoke, Hampshire and New York: Palgrave Macmillan.

Western, Bruce and Deirdre Bloome. 2011. “Cohort Change and Racial Differences in Educational and Income Mobility.” *Social Forces* 90(2): 375–395.

Yum, Minchul. 2015. "Parental Time Investment and Human Capital Formation: A Quantitative Analysis of Intergenerational Mobility." 2015 Meeting Papers 996, Society for Economic Dynamics.

Table 1. Demographic Statistics, Income Sources, and Missing Information (Unweighted Percentages)

Variables	Samples			
	2004-5y	2004-9y	2010-5y	2010-9y
Child's gender (% female)	49.3	49.2	49.5	49.4
Child's age				
29	23.7	23.7		
30	24.2	24.1		
31	25.2	25.2		
32	26.9	27.1		
35			23.7	23.7
36			24.2	24.1
37			25.1	25.1
38			27.0	27.1
Child's income information				
Return	90.5	90.6	88.7	88.8
W-2 + UI	3.6	3.6	4.2	4.1
CPS-based imputation	5.8	5.8	7.1	7.1
Number of missing years of parental information				
0	98.8	95.2	98.8	95.2
1	0.9	2.5	0.9	2.5
2	0.3	1.5	0.3	1.5
3		0.8		0.8
Sample size	12,696	12,608	12,558	12,469

Note:

Children with more than 2 or 3 missing years of parental information are excluded from samples in which sample-selection rules are applied with the five- or the nine-year parental-income measures, respectively

Table 2. Income and Parental-Age Statistics (Weighted Values)

Variables	Samples			
	2004-5y	2004-9y	2010-5y	2010-9y
Child's total income				
Mean	55,121	55,370	69,145	69,329
Standard deviation	75,838	75,677	107,983	107,061
Average parental total income				
Mean	72,097	74,933	72,002	74,826
Standard deviation	107,330	116,052	106,540	115,622
Average parental age				
Mean	43.5	45.3	43.5	45.3
Standard deviation	6.3	6.2	6.3	6.2

Notes:

Children with more than 2 or 3 missing years of parental information are excluded from samples in which sample-selection rules are applied with the five- or the nine-year parental-income measures, respectively. Monetary values in 2010 dollars (adjusted by inflation using the CPI-U-RS).

Table 3. Baseline IGE estimates

	Chetty et al.'s (2014) data	SOI-M best sample	
IGE_e			
Constant	**0.34** (0.32-0.35)[a]	**0.46** (0.43-0.49)[c]	
Nonparametric	**0.38**[b]	**0.50** (0.45-0.54)[c]	
IGE_g		Lower-bound	Upper-bound
Constant	**0.34** (0.34-0.34)[a]	**0.46** (0.40-0.51)[c]	**0.67** (0.59-0.75)[c]
Nonparametric	**0.39**[b]	**0.53** (0.48-0.58)[c]	**0.74** (0.66-0.80)[c]

Notes:

Point estimates are in bold, 95 % confidence intervals are in parentheses. The Online Appendix D provides information on the computation of confidence intervals (and explains why confidence intervals are not available for the nonparametric IGEs based on Chetty et al.'s data). With the SOI-M best sample, lower- and upper-bound estimates of the IGE_g are those in which CPS nonadmins with zero income are dropped and assigned $1, respectively, when computing values for the mean imputation of log income to nonadmin children.

Sources:

[a] Chetty et al. (2014:1574 and Online Appendix C) for the IGE_e and Chetty et al. (2014: Table 1) for the IGE_g; [b] estimated by the authors using data available at https://opportunityinsights.org/data/ ("Geography of Mobility: National Statistics by Parent or Child Income Percentile"); [c] estimated by the authors using the SOI-M best sample.

Table 4. Model-specific Shapley decompositions of total bias in IGE_e estimation

	Constant IGE_e		Nonparametric IGE_e	
	Value	Percent	Value	Percent
Total bias	0.08	100.0	0.09	100.0
Attenuation bias	0.03	42.2	0.03	36.6
Lifecycle bias	0.06	74.8	0.07	79.5
Imputation bias	-0.01	-17.0	-0.01	-16.1

Note:

Total bias is the difference between an estimate based on the SOI-M best sample and the corresponding estimate based on the SOI-M all-biases sample (see Figure 5).

Table 5. Across-models Shapley decomposition of total bias in IGE_e estimation

	Value	Percent
Total bias	0.12	100.0
Functional-form bias	0.03	28.0
Attenuation bias	0.03	28.3
Lifecycle bias	0.06	55.7
Imputation bias	-0.01	-12.1

Note:

Total bias is the difference between the nonparametric estimate based on the best SOI-M sample and the constant-elasticity estimate based on the SOI-M all-biases sample (see Figure 5).

Table 6. Model-specific Shapley decompositions of total bias in IGE_g estimation

Biases with respect to lower-bound estimates				
	Constant IGE_g		Nonparametric IGE_g	
	Value	Percent	Value	Percent
Total bias	0.15	100.0	0.13	100.0
Attenuation bias	0.06	40.7	0.02	18.5
Lifecycle bias	0.04	27.8	0.06	43.7
Selection bias	0.05	31.5	0.05	37.8
Biases with respect to upper-bound estimates				
	Constant IGE_g		Nonparametric IGE_g	
	Value	Percent	Value	Percent
Total bias	0.37	100.0	0.33	100.0
Attenuation bias	0.09	23.4	0.04	11.9
Lifecycle bias	0.01	3.3	0.04	11.6
Selection bias	0.27	73.3	0.26	76.6

Note:
Total bias is the difference between a lower-bound (top panel) or upper-bound (bottom panel) estimate based on the SOI-M best sample and the corresponding estimate based on the SOI-M two-year all-biases sample (see Figure 6).

Table 7. Across-models Shapley decomposition of total bias in IGE_g estimation

Biases with respect to lower-bound estimates		
	Value	Percent
Total bias	0.23	100.0
Functional-form bias	0.09	37.8
Attenuation bias	0.04	18.7
Lifecycle bias	0.05	21.5
Selection bias	0.05	21.9
Biases with respect to upper-bound estimates		
	Value	Percent
Total bias	0.43	100.0
Functional-form bias	0.08	18.5
Attenuation bias	0.06	14.2
Lifecycle bias	0.03	6.2
Selection bias	0.26	61.0

Note:

Total bias is the difference between a lower-bound (top panel) or upper-bound (bottom panel) nonparametric estimate based on the best SOI-M sample and the constant-elasticity estimate based on the SOI-M two-year all-biases sample (see Figure 6).

Figure 1. Four IGEs

		Measure of central tendency	
		Expectation	Geometric mean
Shape of intergenerational curve in log-log space	Straight line	Constant IGE_e	Constant IGE_g
	Smooth curve	Nonparametric IGE_e	Nonparametric IGE_g

Figure 2. Estimators

IGE	Data	
	Chetty et. al (2014a)	SOI-M Panel
Constant IGE_g	*OLS log-log estimator* Estimation of Equation [1] by Ordinary Least Squares	*OLS log-log estimator* Estimation of Equation [1] by Ordinary Least Squares
Constant IGE_e	*Two-step estimator* Estimation of Equation [2] in two steps: 1. Nonparametric estimation of ln E(Y|ln x) at 100 values of X by computing ln mean Y and ln mean X within centile bins of X 2. OLS regression of the estimated ln E(Y|ln x) on ln x	*PPML estimator* Estimation of Equation [2] by Pseudo Maximum Likelihood, using the log-likelihood function of a Poisson regression
Nonparametric IGE_g	*Two-step estimator* 1. Nonparametric estimation of 100 points in the intergenerational curve defined by Equation [7] by computing mean ln Y and ln mean X within centile bins of X 2. Numerical computation of average slope of the curve	*Two-step estimator* 1. Nonparametric estimation of 196 points in the intergenerational curve defined by Equation [7] based on a local polynomial regression of ln Y on ln X 2. Numerical computation of average slope of the curve
Nonparametric IGE_e	*Two-step estimator* 1. Nonparametric estimation of 100 points in the intergenerational curve defined by Equation [6] by computing ln mean Y and ln mean X within centile bins of X 2. Numerical computation of average slope of the curve	*Two-step estimator* 1. Nonparametric estimation of 196 points in the intergenerational curve defined by Equation [6] based on a local polynomial regression of Y on X 2. Numerical computation of average slope of the curve

Note:

To estimate all IGEs, the long-run variables in Equations [1], [2], [6] and [7] are replaced by proxy short-run variables.

Figure 3. Comparison of Chetty et al.'s (2014) data and SOI-M best sample

Year

Cohort	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
1982										14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1981										15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1980										16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
1979																										
1978																										
1977																										
1976																										
1975	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35		
1974	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36		
1973	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37		
1972	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38		

Chetty et al. (2014)

SOI-M Panel

Data used to compute parental income for IGE estimation (preferred estimates)

Data used to compute children's income for IGE estimation (preferred estimates)

Children's ages and years at which parental income is measured

Children's ages and years at which children's income is measured

Figure 4. Chetty et al.'s (2014) children's ages and parental-income measure approximated with data from the SOI-M Panel

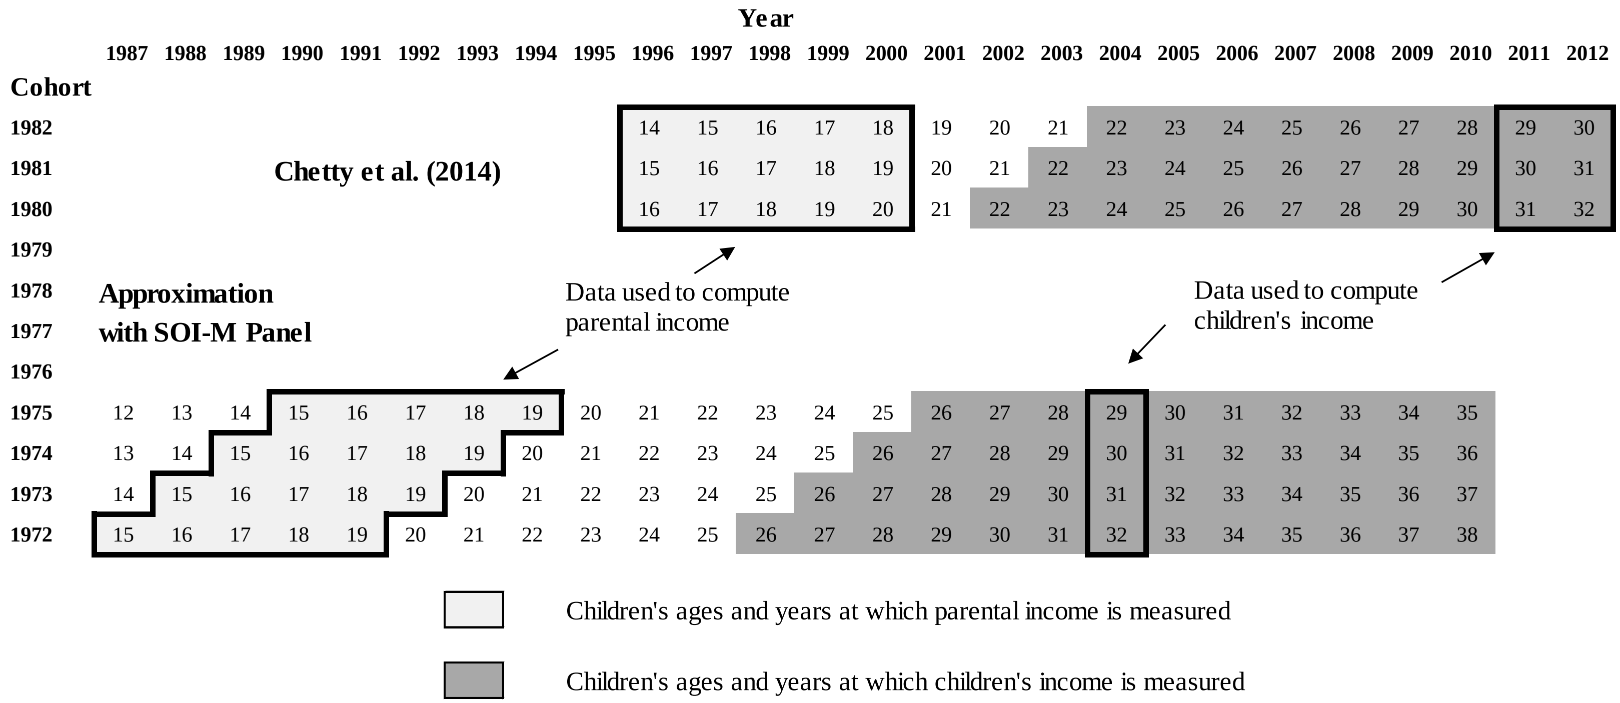

Figure 5. IGE_e estimates based on the SOI-M best sample, the SOI-M all-biases sample, and Chetty et al.'s (2014) data

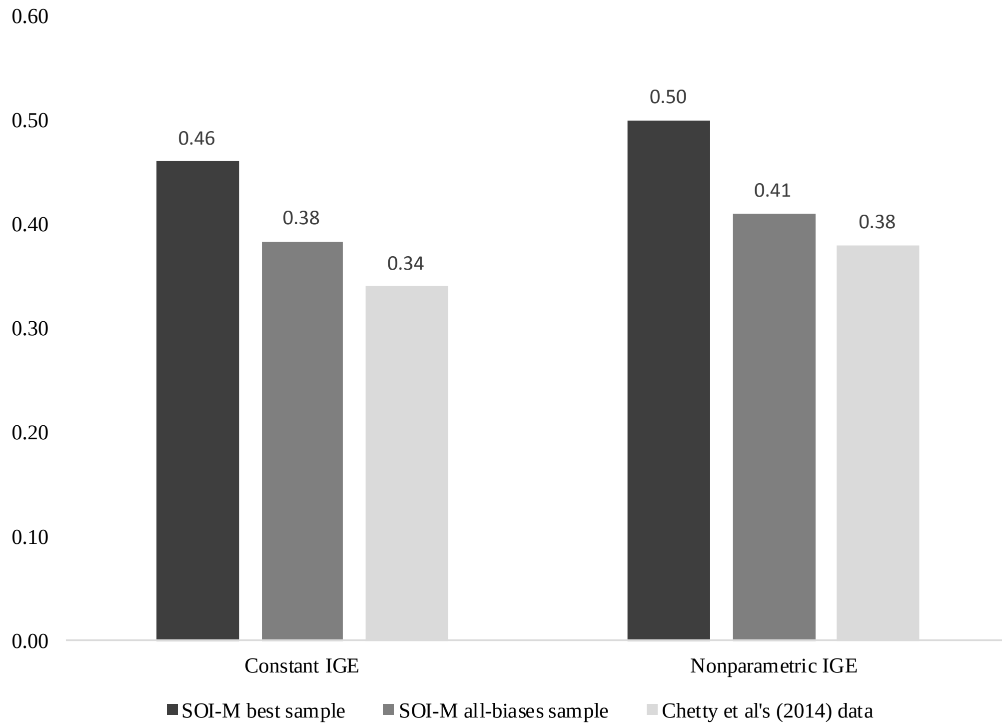

Figure 6. IGE_g estimates based on the SOI-M best sample, the SOI-M all-biases samples, and Chetty et al.'s (2014) data

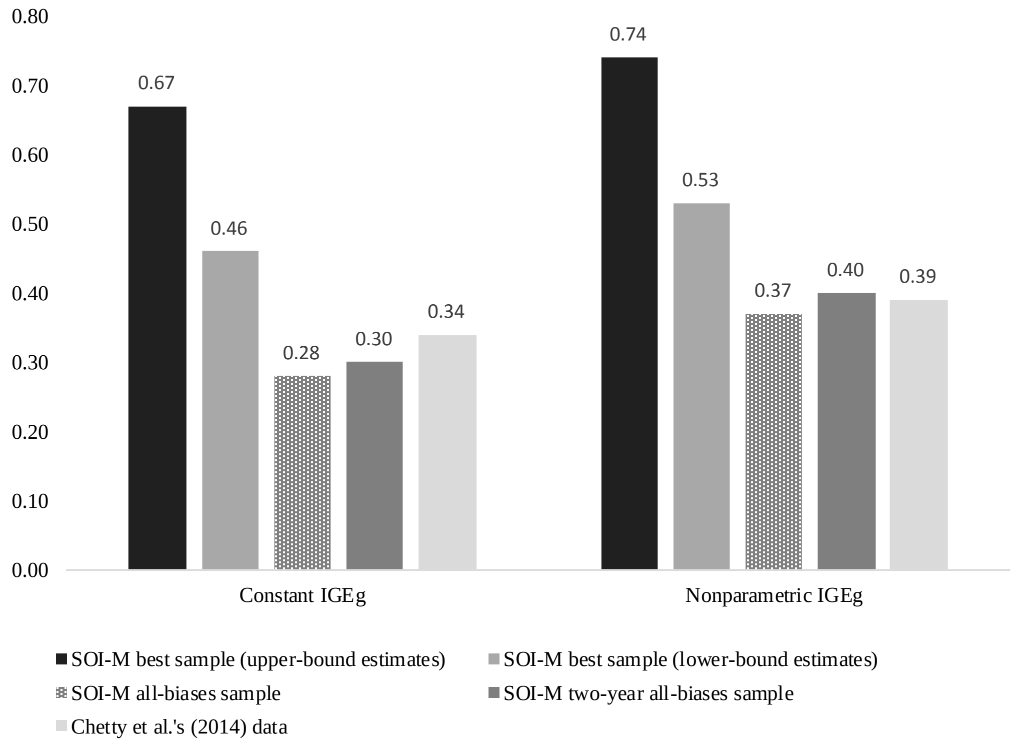

Appendices

THE INTERGENERATIONAL TRANSMISSION OF FAMILY-INCOME ADVANTAGES IN THE UNITED STATES

Pablo A. Mitnik
Stanford Center on Poverty and Inequality
(pmitnik@stanford.edu)

Victoria Bryant
Statistics of Income Division, Internal Revenue Service
(victoria.l.bryant@irs.gov)

Michael Weber
Statistics of Income Division, Internal Revenue Service
(michael.e.weber@irs.gov)

July, 2018

CONTENTS

A. Did Mazumder (2005) overestimate economic persistence?

Chetty et al.'s (2014) criticism of Mazumder's (2005) boils down to the claim that Mazumder's approach for dealing with missing parental data was tantamount to instrumental variable (IV) estimation and therefore led to overly high IGE estimates.[1] Mazumder (2016) has disputed this contention. In particular, he reiterated an argument he had made in his 2005 paper, according to which (a) IV estimates may well be downward biased when both right-side and left-side lifecycle downward biases are present, and (b) this was most likely the case with his data, given that children were "too young" while parents were "too old" at the time their income was measured (Mazumder 2016:92).[2] This argument is important because it entails that Chetty et al.'s contention should be interpreted as a plausible hypothesis that cannot be adjudicated without additional evidence.

Unfortunately, the only additional empirical evidence available on this issue is a robustness analysis reported in Table 6 of Mazumder (2005) that, despite Mazumder's protestation to the contrary (Mazumder 2016:92-93), is quite inconclusive. According to Mazumder (2016), his robustness analysis suggests that his 2005 estimates are not upward biased. In our view, however, the results in that analysis are not inconsistent with the claim that his estimates are upward biased. In particular, it is important in this regard that the IGE estimates in his Table 6 change very little when switching from a seven- to a ten-year measure of parental income once topcoded parents (i.e., those parents for whom imputation was needed) are dropped. Mazumder does provide a possible "selection-effect" explanation for this fact, but it is not clear that this hypothesis should have the upper hand vis-à-vis Chetty et al.'s (2014) hypothesis.

Therefore, our view is that the arguments articulated by Chetty et al. (2014) and Mazumder (2016), respectively, for and against the notion that Mazumder (2005) overestimated persistence are inconclusive.

B. Correct interpretation of the conventionally estimated IGE

Mitnik and Grusky (2017) have shown that the conventionally estimated IGE has been widely misinterpreted: While mobility scholars have assumed that they estimated the elasticity of the *expectation* of children's income, in the general case they actually estimated the elasticity of the *geometric mean* of children's income. That the latter is the case follows immediately from exponentiating and taking natural logarithm on the left-hand side of Equation [1]. Recalling that $GM(W) \equiv \exp E(\ln(W))$, for W any random variable, it is then easy to see that Equation [1] is equivalent to $[1']$. It immediately follows that β_1 is the percentage differential in the geometric mean of children's long-run income with respect to a marginal percentage differential in parental long-run income.

The parameter β_1 is (also) the IGE of the expectation when the population error term satisfies very special conditions (see Santos Silva and Tenreyro 2006; Petersen 2017; Wooldridge 2002:17), but not otherwise. The parameter $\beta_1$would (also) be the elasticity of the conditional expectation of the children's income in the general case if it were the case that $E(\ln Y|x) = \ln E(Y|x)$, but this equality does not hold (due to Jensen's inequality).

C. Share interpretation of nonparametric IGEs

In the main text we argued that the nonparametric IGE_e provides an approximation to the ratio between the inequality in children's opportunities and the inequality in parental income (the quantity in the right-hand side of Equation [5]). We show here why that is the case.

From Equation [6]:

$$SD(\ln E(Y|X)) = \mathrm{SD}\big(\mathrm{F}(\ln \mathrm{X})\big).$$

A first-order Taylor-series approximation to the variance of the random function $F(.)$ centered around $E(\ln X)$ gives:

$$\mathrm{Var}\big(\mathrm{F}(\ln \mathrm{X})\big) \approx \left(\frac{\partial F(t)}{\partial t}\bigg|_{\mathrm{t = E(ln\ X)}} \right)^2 Var(\ln X),$$

so

$$\frac{\mathrm{SD}\big(\mathrm{F}(\ln \mathrm{X})\big)}{SD(\ln X)} \approx \frac{\partial F(t)}{\partial t}\bigg|_{\mathrm{t = E(ln\ X)}}.$$

Intergenerational curves are well approached by third-degree polynomials. We may then write:

$$\ln E(Y|x) = F(\ln x) \approx \gamma_0 + \gamma_1 \ln x + \gamma_2 (\ln x)^2 + \gamma_3 (\ln x)^3.$$

We then have:

$$\frac{\partial F(t)}{\partial t}\bigg|_{\mathrm{t = E(ln\ X)}} \approx \gamma_1 + 2\,\gamma_2 E(\ln X) + 3\,\gamma_3 [E(\ln X)]^2. \qquad [C1]$$

Let's denote the nonparametric $\mathrm{IGE_e}$ by $\bar{\eta}$. We then have:

$$\bar{\eta} = E\left(\frac{\partial F(\ln X)}{\partial \ln X}\right) \approx E\left(\frac{\partial[\gamma_0 + \gamma_1 \ln X + \gamma_2 (\ln X)^2 + \gamma_3 (\ln X)^3]}{\partial \ln X}\right)$$

$$\approx \gamma_1 + 2\,\gamma_2 E(\ln X) + 3\,\gamma_3 E((\ln X)^2). \qquad [C2]$$

It follows from Equations [C1] and [C2] that

$$\frac{\partial F(t)}{\partial t}\bigg|_{\mathrm{t = E(ln\ X)}} \approx \bar{\eta} - 3\,\gamma_3 Var(\ln X).$$

As, empirically, γ_3 is very small (if not statistically indistinguishable from zero), $\bar{\eta}$ provides an approximation to $\frac{\partial F(t)}{\partial t}\Big|_{\mathrm{t = E(ln\ X)}}$ and therefore to $\frac{SD(\ln E(Y|X))}{SD(\ln X)}$. Of course, the closer to quadratic the intergenerational curve (i.e., the closer to zero γ_3 is), the better this approximation will be (and vice versa). With our data, $\frac{SD(\ln E(Y|X))}{SD(\ln X)}$ is approximately six percent smaller than $\bar{\eta}$.

A similar analysis applies in the case of the nonparametric IGE_g.

D. Estimation and statistical inference

We provide here a detailed discussion of the estimators employed to estimate the four IGEs in Figure 1, which were briefly described in the main text. We also provide details on the computation and interpretation of the confidence intervals reported in the main text, and explain why we did not follow the customary approach of including polynomials on children's and parents' age as control variables when estimating IGEs.

Estimators

Figure 2, in the main text, provides a summary of our discussion in this section.

The estimates of the constant IGE_g are all the result of estimating the PRF of Equation [1] by OLS, i.e., of employing the "OLS log-log estimator." This applies to the estimate reported by Chetty et al. (2014) and to the estimates based on the SOI-M data we obtained.

The estimates of the constant IGE_e based on the SOI-M data we produced were generated by estimating Equation [2] with the PPML estimator (Santos Silva and Tenreyro 2006). Like the OLS log-log estimator, the PPML estimator is semiparametric, that is, it makes no assumption regarding the distribution of the dependent variable and is consistent as long as the mean function is correctly specified (Gourieroux, Monfort, and Trognon 1984).[3]

The estimate of the constant IGE_e reported by Chetty et al. (2014) is based on a two-step estimator of Equation [2]. In the first step, nonparametric estimates of $\ln E(Y|\ln x)$ are generated by binning the children's parental income into 100 equal-sized (centile) bins, computing the mean income of parents and children within each bin, and taking the natural logarithm of those means. In the second step, an estimate of α_1 is generated by running an OLS regression of the estimates of $\ln E(Y|\ln x)$ on the corresponding $\ln x$ values.

As indicated in the main text, the estimators of the nonparametric IGEs are all two-step estimators: The first step produces nonparametric estimates of a number of points in the relevant intergenerational curve—i.e., the curve defined by either Equation [6] or Equation [7]—while the second step estimates the average slope of the curve through a numerical approximation based on the estimated points. Across datasets, the estimators only differ on the nonparametric approach used to estimate the points of the intergenerational curves and on the number of points that are estimated and employed in the numerical approximations.

The estimates of the nonparametric IGEs based on Chetty et al.'s (2014) data we produced for this article rely on estimates of the points of the intergenerational curves they obtained by (a) binning the children's parental income into centile bins, and (b) computing the log mean income of parents and the log mean income and mean log income of children within each bin.[4] These points pertain, approximately, to the 0.5, 1.5, …, 98.5, 99.5 percentiles of parental income. The estimates based on the SOI-M data we produced rely on estimates of 196 points pertaining to the 1,1.5, 2, . . ., 98, 98.5, 99 percentiles of parental income. We obtained these estimates by resorting to local polynomial regressions (Cleveland, Devlin, and Grosse 1988; Cleveland and Grosse 1991) of $\ln Y$ on $\ln X$ (in the case of the nonparametric IGE_{g}) and of Y on X (in the case of the nonparametric IGE_{e}). In all cases we used degree 1 polynomials and a tricube weight function. The smoothing parameter, however, is specific to each estimated curve, as it is the global minimizer of the AICc—a version of the Akaike Information Criterion specifically tailored to nonparametric regression (Hurvich, Simonoff, and Tsai 1998)—within the range of [.08, 1].

Regardless of dataset and IGE concept, in the second step the nonparametric IGE is estimated as follows. First, the expected point elasticity (or average slope of the curve) between any two contiguous points estimated in the first step is approximated by computing the corresponding ratio of finite differences. For instance, in the case of the nonparametric IGE_{e}, for two contiguous values of parental income x_1 and x_2, the approximate average elasticity between those points is computed by dividing $(\ln E(Y|x_2) - \ln E(Y|x_1))$ by $(\ln x_2 - \ln x_1)$ (where all quantities refer to estimates generated in the first step). Second, the nonparametric IGE estimate is computed as a simple average of the estimated elasticities between all pairs of contiguous points. Importantly, all segments defined by contiguous points cover the same share of children in the population: 0.5 percent with the points estimated with the SOI-M data and about 1 percent with the points estimated with Chetty et al.'s (2014) data.[5] Therefore, the final average is based on a set of self-weighting partial averages.[6]

Computation and interpretation of confidence intervals

In Table 3 we report 95 percent confidence intervals for the IGE estimates. In the case of the IGE estimates from Chetty et al. (2014), we simply report the confidence intervals implicit in the provided standard errors. In the case of the new estimates of constant IGEs based on the SOI-M data we obtained, we construct confidence intervals based on robust standard errors (which is mandatory with a PML estimator); the standard errors also take into account the clustering of children into families (see, e.g., Rogers 1993).

In the case of the estimates of nonparametric IGEs based on the SOI-M data we produced, statistical inference is based on the nonparametric bootstrap with 2,000 bootstrap samples. In the nonparametric context, use of the bootstrap-based percentile method (Efron and Tibshirani 1986) produces "variability bounds" (Racine 2008) or "confidence bands" (Wasserman 2006), which can be interpreted as approximations to true confidence intervals. For simplicity of terminology, in Table 3 we refer to the resulting intervals as "confidence intervals."

We cannot compute similar (approximate) confidence intervals for the nonparametric IGE estimates based on Chetty et al.'s (2014) data we report, as we do not have access to the underlying microdata (and therefore cannot generate the needed bootstrap results).

Age controls

The relationship between long-run income measures (for instance, lifetime average income) and short-run proxy measures varies with the age at measurement (both for parents and children). For this reason, when estimating the constant IGE_g with proxy measures, it has been customary to include polynomials on children's and parents' age as control variables, each indexing the age at which the income measurements were taken. However, Chetty et al.'s (2014) estimates of constant IGEs and our estimates of constant and nonparametric IGEs were produced without including such controls. In both cases, the variability in children's age in the samples is very minor, so controlling for age is unnecessary. In addition, because the age at which parents have their children is not exogenous to their income and the former's age may affect the latter's life chances, Mitnik et al. (2018: Online Appendix K) argued that controlling for parental age is inconsistent with the objective of measuring the gross association between parents' and children's income. In line with this argument and with the comparative purposes of this article, all new estimates we produced were generated without including age controls.

E. Chetty et al.'s (2014) evidence on lifecycle bias: Critical examination

Chetty et al. (2014) claimed that both the (constant) IGE_g and the (constant) IGE_e stabilize around age 30 and that this entails that their IGE estimates are free of left-side lifecycle bias.[7] However, their evidence for this claim is very problematic. As also argued by Mazumder (2016:113-14), in order to estimate the (constant) IGE_g for children older than 32, Chetty et al. resorted to auxiliary samples covering cohorts born earlier than those in their core sample. As the parental income of the children from these auxiliary samples is measured in 1996-2000 (i.e., the same period employed to measure parental income in their core sample), when children's age increases the age at which parental income is measured increases as well (e.g., for children age 41, parents' income is measured when parents are about ten years older than for children age 31). It is well-known that measuring parents' income when they are in their fifties depresses estimates (e.g., Grawe 2006; Haider and Solon 2006; Mazumder 2001). This means that, even if there is a reduction in left-side lifecycle bias as children get older than 32, in Chetty et al.'s analysis this would tend to be masked by the effect of simultaneously increasing the age at which parental income is measured (as the parents of core-sample children are already in their mid-40s in 1996-2000).[8]

On its part, their evidence for the (constant) IGE_e, which is based on the core sample only, does not include estimates for children older than 32. In this case Chetty et al. reported that estimates increase at a decreasing rate as children move from ages 22 to 32, and based on the fact that the estimated IGE is 2.1 percent higher at age 32 than at age 31, they concluded that it stabilizes around age 30. However, even with a decreasing growth rate, a 2.1 percent increase in one year (of age) is not necessarily that small. For instance, if the rate of growth each year is 0.9 of what it was in the previous year (and given their estimate of 0.343 by age 32), we should expect an IGE_e of about 0.38 by age 40, or close to 20 percent higher than by age 30. More generally, the "argument by extrapolation" they offer is not very persuasive, as the observed IGE-age points do not constrain that much the ways we may reasonably image the IGE-age curve continues past age 32. Chetty et al.'s argument is also inconsistent with the lifecycle-bias analyses of Mitnik et al. (2018: Online Appendix I) and Mitnik (2017a), which suggest that estimates of the IGE_e of total family income do not stabilize around age 30.[9]

F. Attenuation bias and the advantage of tax income data

Chetty et al. (2014) claimed that income IGE estimates based on tax data should be less affected by attenuation bias than previously reported in the literature. They provided three arguments: (a) family income fluctuates less than individual earnings across years, (b) income is measured with less error in tax data than in survey data, and (c) the approach Mazumder (2005) employed to deal with missing parental information led him to overestimate the magnitude of attenuation bias. We consider these arguments in turn.

It is generally accepted that family income fluctuates less over time than father's earnings (e.g., Mazumder 2005:250), so the first argument seems unproblematic. On its part, the second argument can be strengthened considerably. The reason is that it's not really necessary that tax data be measured with less error than survey data for attenuation bias to be less of a problem with the former data. As tax data cover much better than survey data the upper tail of the parental-income distribution, which makes the "signal" larger, it should be enough that tax data do not include more "noise" than survey data (Mitnik 2017a:29-30).

This is easiest to see in the case of the constant IGE_{g}. Assuming no lifecycle bias, the standard analysis of attenuation bias in the OLS estimation of this elasticity (e.g., Solon 1992), is that $plim\,\widehat{\beta_1} = \beta_1 \frac{Var(\ln S)}{Var(\ln S)+Var(\varepsilon)}$, where ε is a zero-expectation additive noise in the logarithm of the short-run measure of parental income S with respect to the logarithm of the long-run measure X. This entails that even if $Var(\varepsilon)$ were the same with survey and tax data, the "attenuation factor" multiplying β_1 would still be closer to one with tax data due to their better coverage of the upper tail of the parental-income distribution, as this can be expected to lead to a larger value of $Var(\ln S)$. Mitnik (2017a) has shown that a similar analysis applies in the case of the constant IGE_{e}.

The third argument is, of course, closely related to the argument we considered in our Online Appendix A. There is, however, a subtle but consequential difference between the two arguments. In that appendix, the focus was on the estimates Mazumder (2005) produced with the measure of father's earnings based on the maximum number of years of information he had available, with Chetty et al. (2014) identifying a feature of the data (imputation of father's earnings) that can be expected to push the estimates up and Mazumder (2016) identifying a different feature of the data (children's and fathers' ages) that can be expected to push them down. In contrast, here the focus is on the full series of estimates Mazumder obtained when increasing the years of information used to compute the measure of father's earnings. As Mazumder (2005) imputes earnings for a larger share of fathers when he uses more years of information, the estimates increase not only because the variance of the error $Var(\varepsilon)$ falls but also because the estimates are moving toward IV estimates. In this context (a) the children's ages remain constant over the analysis, and (b) the fathers' ages decrease when more years of information are used, which in this case can be expected to reduce *downward* right-side lifecycle bias (as the average age of parents is 47 with the parental measure based on the fewest years of information). Therefore, it is quite likely that the estimates of attenuation bias computed from Mazumder's results are upward biased.[10]

We therefore conclude that Chetty et al.'s (2014) claim that, for estimates based on tax income data, attenuation bias should be a less serious problem than previously reported in the literature has clear merit.

G. Chetty et al.'s (2014) evidence on attenuation bias: Critical examination

Chetty et al. (2014) claimed that their estimates of the constant IGE_g based on tax data nearly stabilize once five years are employed, and that this entails that they are (nearly) free of attenuation bias. To support their claim, they reported that the constant IGE_g of family income hardly increased (i.e., from 0.344 to 0.366, or six percent) when they used 15 years of parental information instead of five.[11]

Mazumder (2016:114-115) has forcefully criticized this evidence. He pointed out that when Chetty et al. constructed their measure of parental income based on 15 years of information, all of the additional years pertained to when the parents were older. For instance, if for a specific child the five-year measure of parental income pertains to when her parents are 42-46 years old (in 1996-2000), the fifteen-year measure pertains to when they are 42-56 years old (in 1996-2010). But, as we mentioned in Online Appendix E, measuring parents' income when they are old depresses estimates; it is for this reason that it has long been argued that attenuation bias is best reduced by adding parental information from parents' prime-age period, not by adding information when they are in their fifties (Mazumder 2005:247-248). Moreover, Mazumder (2016:114-115) has provided survey-data evidence consistent with his argument, which suggest that adding years of parental information "in the wrong direction" may even *reduce* estimates if the added years pertain to when the parents are old enough. Mazumder concluded that the results that Chetty et al. (2014) reported were most likely distorted by the increasing noisiness of the additional years of parental information they used in their analysis.

Chetty et al. did consider this possibility, but rejected it on the argument that they had provided evidence that "estimates of mobility are not sensitive to varying the age in which parent income is measured over the range observed in our dataset" (2014: Online Appendix E, ftn. 9). The evidence in question, however, pertains to the rank-rank slope. The rank of parents may remain the same as they get older even as the differences between their incomes increase. As the latter can be expected to raise IGEs, the argument does not hold much water.

In Mitnik et al.'s (2018: Online Appendix H) gender-specific attenuation-bias analyses with tax data, they focused on the IGE_e rather than the IGE_g. They computed the constant IGE_e with parental income measures based on one to nine years of information, and reached three conclusions: (a) attenuation bias is greatly reduced by using nine years of parental information, (b) using five years instead of nine years would result in a non-negligible increase in bias, and (c) although estimates appear to be reaching a plateau once nine years of parental information are employed, it is not possible to rule out that some (downward) bias remains. Mitnik et al.'s evidence is therefore inconsistent with the notion that, with tax-based data, five years of information are enough to eliminate the bulk of attenuation bias.

H. Values used in the mean imputation of income and log income for nonadmin children

The Annual Social and Economic Supplement of the Current Population Survey (CPS-ASEC) identifies likely nonfilers using a tax simulation model. Although this information is available for the entire period covered by the SOI-M Panel, the CPS-ASEC data after 2003 have serious inconsistencies and cannot be used. We therefore use pooled CPS-ASEC data from 1999 to 2003 to estimate the mean income of nonadmins by gender-age group. The resulting mean values, which we used for income imputation, are as follows (all in 2010 dollars): 26-30 year-old men: $4,910; 31-35 year-old men: $5,815; 36-40 year-old men: $6,706; 26-30 year-old women: $5,372; 31-35 year-old women: $6,574; 36-40 year-old women: $7,560.

Approximately one-third of CPS nonadmins have zero family income. Therefore, simply computing their mean log income by gender-age group is not feasible. As discussed in the main text, we computed two sets of values. While in one case CPS nonadmins with zero income are dropped, in the other they are assigned an income of $1. The resulting mean values are the following (the first figure in each pair pertains to the computation in which CPS nonadmins with zero income are dropped): 26-30 year-old men: 8.25 and 3.84; 31-35 year-old men: 8.40 and 4.72; 36-40 year-old men: 8.62 and 5.33; 26-30 year-old women: 8.42 and 4.80; 31-35 year-old women: 8.49 and 5.67; 36-40 year-old women: 8.63 and 6.04. These mean log income figures provide upper and lower imputation values leading, respectively, to lower and upper estimates of the IGE_g.

I. Additional attenuation-bias results

Given Chetty et al.'s (2014) strong denial that their IGE_g estimates are affected by attenuation bias, it seems imperative that we examine whether our empirical findings with regards to attenuation bias are robust to the various ways in which attenuation-bias analyzes can be conducted. This is what we do here, focusing on the constant IGEs (as in the previous literature on attenuation bias) for men and women pooled. We also provide and discuss evidence relevant for assessing whether the estimates based on the SOI-M data we produced using nine-year parental-income measures are likely to be (nearly) free of attenuation bias.

In Table I1 we present estimates of the constant IGE_e and IGE_g obtained with parental measures based on five, eight and nine years of information. We generated these estimates with children's income measured either in 2004 or 2010—that is, when the children were in their early or late thirties—and using either the "common-rules" or the "common-sample" approach. Briefly, the common-rules approach uses (nearly) the same sample inclusion rules regardless of the number of years of information employed to compute the parental income measure; in contrast, the common-sample approach uses (nearly) the same sample to generate all estimates, i.e., the sample that is selected when the inclusion rules are applied with a particular *n*-year parental variable, regardless of the number of years of information actually employed to compute the parental variable used for estimation (for details, see Mitnik et al. 2018: Online Appendix H).

Here we implemented this approach for $n = 9$ (common-sample approach I) and $n = 5$ (common-sample approach II).[12] In addition, in the case of the IGE_g, we generated the estimates shown in Table I1 after assigning nonadmin children $0 income (i.e., dropping them from the analysis, as in Chetty et al. 2014), or CPS-based mean log income values (computed after dropping CPS nonadmins with zero income, as in the lower-bound estimates in the second column of Table 3).

The results for the three approaches are displayed in Table I1 in contiguous horizontal panels, each of which includes five columns. We start by focusing on the first and third columns of each panel—which show IGE estimates based on five- and nine-year parental-income variables—and on the fourth column—which shows the percent difference between these two estimates. The results in these columns are uniformly inconsistent with Chetty et al.'s (2014) contention that five years of parental information are enough to nearly eliminate attenuation bias. Regardless of approach, IGE concept, children's ages, and treatment of nonadmin children, the estimates increase substantially—between 7.5 and 9.1 percent for the IGE_e, and between 15.7 and 33.0 percent for the IGE_g—when switching from the five-year to the nine-year parental measure. Therefore, we can conclude that our evidence that IGE estimates based on five-year parental measures are affected by substantial attenuation bias is very robust.

The second and third columns in each panel allow to compare IGE estimates based on eight- and nine-year parental-income variables, while the fifth column shows the corresponding percent differences. The comparison of estimates based on eight and nine years of information provides some evidence on whether the latter are likely to be (nearly) free of attenuation bias; if this is the case, we should observe that the differences between estimates are very small. The results on this regard differ across IGE concepts, and also across treatments of nonadmin children when estimating the IGE_g.

In the case of the IGE_e, no difference in estimates is larger than one percent and the differences are half of that or less in the case of the common-rule approach. This is exactly what we would expect if the nine-year estimates had (almost) converged to the long-run estimates of interest, so these findings are a reason for optimism. At the same time, as Mitnik et al. (2018: Online Appendi H) have also pointed out, evidence like this does not allow to rule out that a non-negligible amount of attenuation bias still remains.[13]

In the case of the IGE_g with CPS-based mean imputation for nonadmin children, estimates increase between 1.9 and 5.1 percent when we use nine instead of eight years of parental information. When nonadmin children are instead assigned $0 and therefore dropped from the analyses, estimates increase significantly less, between 1.5 and 2.2 percent—but still significantly more than in the case of the IGE_e. It therefore seems safe to conclude that, in both cases, the IGE_g estimates based on nine years of parental information are affected by residual attenuation biases of a non-negligible magnitude. This provides further support for our interpretation of the IGE_g estimates obtained by dropping CPS nonfilers with zero income (when computing values for mean imputation) as *lower-bound* estimates, and therefore for our claim that Chetty et al.'s (2014) IGE_g estimates greatly understated true economic persistence.

J. Shapley decompositions of total bias in IGE_g estimation

In the main text we indicated that to properly capture the effects of lifecycle, attenuation, functional-form and selection bias on the estimates based on Chetty et al. (2014) data and methodological decisions, all estimates should be obtained with two-year samples, but that proceeding this way was not feasible. Here, we explain why this is the case as well as the alternative approach we used.

The first problem is that, with two-year samples, (a) the relevant mean values for imputation are not the mean annual log income values employed to produce the estimates shown in Table 3, due to the dependencies between children's income across years, and (b) the auxiliary data that would be needed to compute mean values that take into account those dependencies are not available (the CPS-ASEC data are unsuitable for this purpose). The second problem is that some two-year measures would require using 2009 information to compute them. However, as reported by Mitnik et al. (2018: Online Appendix I), income data for 2008 and 2009 appear to be seriously affected by the Great Recession, which greatly compressed the income distribution. Using the 2009 data to compute income IGEs is therefore unadvisable.

As an alternative, we proceeded as follows. As we did in the case of the IGE_e, we computed 16 estimates, using in all cases one-year samples. However, before computing the Shapley decompositions, we adjusted all estimates based on samples where nonadmin children's income was assumed to be zero. We did this adjustment separately for estimates of the constant and the nonparametric IGE_g, under the assumption that they all underestimate the IGE_g by the same amounts that the corresponding estimates based on the SOI-M (one-year) all-biases sample underestimate it (compared to the estimates based on the SOI-M two-year all-biases sample; see Figure 6). After introducing these adjustments, we computed the Shapley decompositions as before.

Importantly, with this approach, the total biases are exactly what is wanted in all cases. The total biases we would like to decompose are in all cases differences between the IGE_g estimates based on the SOI-M best sample and the SOI-M two-year all-biases sample. As the *adjusted* estimates based on the SOI-M (one-year) all-biases sample are identical to the estimates based on the SOI-M two-year all-biases sample, the total biases we actually decompose (differences between the estimates based on the SOI-M best sample and the *adjusted* estimates based on the SOI-M all-biases sample) are equal to the total biases of interest. We make use of this equivalence in defining total biases in the notes to Table 6 and 7.

Notes

[1] For the reason why IV estimates are usually expected to be upward biased in the intergenerational-mobility context, see Solon (1992: Appendix) and Mitnik (2017b:8-10).

[2] See Mitnik (2017b: Eqs. 13 and 14) for a measurement-error model consistent with this argument.

[3] See Mitnik (2017c) for how to estimate the constant IGE_e using the PPML estimator and the statistical package Stata.

[4] The estimated points have been made publicly available by Chetty et al. (2014); see the sources in Table 3 for details.

[5] The latter under the assumption that, within centiles of parental income, mean and median parental income are approximately the same.

[6] In the case of estimates based on the SOI-M data, the estimation of nonparametric IGEs ignores the intergenerational curves' final left and right segments (each covering 1 percent of children). Because the curve is estimated less precisely at the boundaries, these "trimmed estimators" are often more efficient. The point estimates from the trimmed and untrimmed estimators are, however, very similar. With Chetty et al.'s (2014) data, the estimation of the nonparametric IGEs also ignores the final left and right segments (each covering approximately 0.5 percent of children).

[7] For the IGE_g, see Chetty et al. (2014:1580 and Online Appendix Figure IIa). For the IGE_e, see Chetty et al. (2014: Online Appendix C and Figure Ib). Chetty et al. (2014) also claimed that their rank-rank slope estimates are unaffected by lifecycle bias. Our discussion here is only concerned with, and is only meant to apply to, lifecycle bias in the estimation of IGEs.

[8] An additional potential issue is that, in producing their evidence on lifecycle-bias, Chetty et al. (2014) estimated the IGE_g not with the full sample (as is standard), but only with children with parental income between the 10th and the 90th percentiles (see notes to their Online Appendix Figure IIa). Proceeding this way could have had an impact on their results.

[9] Mitnik et al.'s (2018: Online Appendix I) results are not conclusive, as some of their IGE estimates appear to have been driven downward by the Great Recession. Nevertheless, overall they do suggest that IGE estimates based on a sample of men and women that are 29-32 years old when their income is measured should be affected by lifecycle bias.

[10] This could still not be the case if, for instance, negative lifecycle biases are substantially larger in absolute value with the IV than with the OLS estimator. There is no evidence, however, that this is so.

[11] See Chetty et al. (2014:Table 1 and Online Appendix E). Chetty et al. (2014) also claimed that their rank-rank slope estimates are unaffected by attenuation bias. Our discussion here is only concerned with, and is only meant to apply to, attenuation bias in the estimation of IGEs.

[12] In our Shapley decompositions we used the common-rules approach. The attenuation-bias evidence provided by Chetty et al. (2014; see our Online Appendix G) is likely based on the common-sample approach (with $n = 5$). Mitnik et al. (2018: Online Appendix H) and Mazumder (2005) used both approaches in their attenuation-bias analyses.

[13] First, the evidence in Table I1 is not inconsistent with the estimates converging very slowly, past year nine, to the long-run estimates; if this were the case, the long-run estimates of interest could still be significantly larger than those based on nine years. Second, the argument advanced by Mazumder (2016) in his criticism of Chetty et al.'s (2014) attenuation-bias evidence is relevant here as well: It is possible that the nine-year estimates only appear to have (almost)

converged to the long-run estimates of interest because the years of income information that are being added as we move from the five-year measure to the nine-year measure pertain to the “wrong parental ages” (see our Online Appendix G).

References

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

Cleveland, William, Susan Devlin, and Eric Grosse. 1988. "Regression by Local Fitting: Methods, Properties, and Computational Algorithms." *Journal of Econometrics* 37(1): 87-114.

Cleveland, William and Eric Grosse. 1991. "Computational Methods for Local Regression." *Statistics and Computing* 1: 47-62.

Efron, Bradley and Robert Tibshirani. 1986. "Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy." *Statistical Science* 1(1): 1-154.

Gourieroux, C, A. Monfort and A. Trognon. 1984. "Pseudo Maximum Likelihood Methods: Theory." *Econometrica* 52(3): 681-700.

Grawe, Nathan. 2006. "Lifecycle Bias in Estimates of Intergenerational Earnings Persistence." *Labour Economics* 13: 551-570.

Haider, Steven and Gary Solon. 2006. "Life-Cycle Variation in the Association between Current and Lifetime Earnings." *American Economic Review* 96(4): 1308-1320.

Hurvich, Clifford, Jeffrey Simonoff, and Chih-Ling Tsai. 1998. "Smoothing Parameter Selection in Nonparametric Regression Using an Improved Akaike Information Criterion." *Journal of the Royal Statistical Society* B, 60: 271–293.

Mazumder, Bhashkar. 2001. "The Miss-measurement of Permanent Earnings: New Evidence from Social Security Earnings Data." Federal Reserve Bank of Chicago Working Paper 2001-24.

Mazumder, Bhashkar. 2005. "Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data." *The Review of Economics and Statistics* 87(2): 235-255.

Mazumder, Bhashkar. 2016. "Estimating the Intergenerational Elasticity and Rank Association in the United States: Overcoming the Current Limitation of Tax Data." In *Inequality: Causes and Consequences*, edited by Lorenzo Cappellari, Solomon Polacheck, and Konstantinos Tatsiramos. Bingley: Emerald.

Mitnik, Pablo. 2017a. "Estimating the Intergenerational Elasticity of Expected Income with Short-Run Income Measures: A Generalized Error-in-Variables Model." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo. 2017b. "Intergenerational Income Elasticities, Instrumental Variable Estimation, and Bracketing Strategies." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo. 2017c. "Estimators of the Intergenerational Elasticity of Expected Income: A Tutorial." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo, Victoria Bryant, Michael Weber and David Grusky. 2018. "A Very Uneven Playing Field: Economic Mobility in the United States." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo and David Grusky. 2017. The Intergenerational Elasticity of What? The Case for Redefining the Workhorse Measure of Economic Mobility." Stanford Center on Poverty and Inequality Working Paper.

Racine, Jeffrey. 2008. "Nonparametric Econometrics: A Primer." *Foundations and Trends in Econometrics* 3(1): 1-88.

Rogers, Williams. 1993. "Regression Standard Errors in Clustered Samples." *Stata Technical Bulletin* 13: 19-23.

Santos Silva, J. M. C. and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 88(4): 641-658.

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

Petersen, Trond. 2017. "Multiplicative Models for Continuous Dependent Variables: Estimation on Unlogged versus Logged Form." *Sociological Methodology* 47:113-164.

Wasserman, Larry. 2006. *All of Nonparametric Statistics*. New York: Springer.

Wooldridge, Jeffrey. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass: The MIT Press.

Table I1. IGE estimates with five, eight and nine years of parental information

	Common-sample approach I					Common-sample approach II					Common-rules approach				
	Years of par. inf.			% Δ 5 to 9 years	% Δ 8 to 9 years	Years of par. inf.			% Δ 5 to 9 years	% Δ 8 to 9 years	Years of par. inf.			% Δ 5 to 9 years	% Δ 8 to 9 years
	5	8	9			5	8	9			5	8	9		
Constant IGE_e, CPS-based mean imputation															
Year 2010, children in their later 30s	0.429	0.457	0.461	7.5	0.8	0.427	0.453	0.456	7.0	0.8	0.428	0.460	0.461	7.7	0.3
Year 2004, children in their early 30s	0.368	0.397	0.401	9.1	1.0	0.371	0.396	0.400	7.8	0.8	0.371	0.399	0.401	8.1	0.5
Constant IGE_g, $0 imputation															
Year 2010, children in their later 30s	0.308	0.387	0.393	27.6	1.7	0.318	0.387	0.396	24.5	2.2	0.321	0.386	0.393	22.6	1.8
Year 2004, children in their early 30s	0.271	0.317	0.324	19.6	2.0	0.277	0.323	0.329	18.5	1.8	0.279	0.319	0.324	16.1	1.5
Constant IGE_g, CPS-mean imputation (CPS nonadmins dropped)															
Year 2010, children in their later 30s	0.343	0.434	0.456	33.0	5.1	0.383	0.459	0.470	22.7	2.4	0.385	0.435	0.456	18.5	4.9
Year 2004, children in their early 30s	0.324	0.397	0.416	28.5	4.7	0.359	0.417	0.425	18.5	1.9	0.360	0.399	0.416	15.7	4.3

Notes:

The common-sample approach keeps the sample (nearly) fixed when estimating IGEs using five and nine years of parental information. There are two ways of implementing this approach. In the columns under "Common-sample I" the sample inclusion rules are applied with the nine-year parental variable. In the columns under "Common-sample II" the sample inclusion rules are applied with the five-year variable. The common rules approach keeps the sample inclusion rules (nearly) fixed, and applies them with the five- or the nine-year parental variable, as relevant. "% Δ" denotes "percent difference."