Online Appendices

REJOINDER: A FORCED CRITIQUE OF THE INTERGENERATIONAL ELASTICITY OF THE CONDITIONAL EXPECTATION

Pablo A. Mitnik
Center on Poverty and Inequality
Stanford University

David B. Grusky
Center on Poverty and Inequality
Stanford University

Sociological Methodology, Vol. 50, Issue 1, 2020.

CONTENTS

A. Additional issues raised by Lundberg and Stewart's comment

Due to space constraints, many of our reactions to LS's comment have to be relegated to this appendix. We discuss here three additional issues raised by their comment.

Visual representations of quantile intergenerational curves and scale invariance

We pointed out in Section 5 of our rejoinder that, because of measurement error, estimates of LS's quantile intergenerational curves will be (asymptotically) biased. In this section, we now assume that (a) LS or others have solved the methodological problem of producing consistent nonparametric estimates of long-run quantile intergenerational curves with short-run income measures, or (b) LS or others have sidestepped this problem by claiming that, unlike most of the existing mobility literature, they are not interested in the relationship between families' long-run income and their children's long-run income.

The latter assumption might be the most attractive to LS because it would make their visualizations immediately usable without any additional analytic or empirical methodological work. For instance, LS might suggest that the mobility field's usual interest in long-run income is simply too ambitious and that it should only be pursued retrospectively, once complete income histories are collected in survey panels or administrative registers. By this logic, the field would be well advised to pursue more modest goals, at least until such time as the requisite data are available. If this argument were accepted, the visual representation of quantile intergenerational curves proposed by LS would not be subject to our main methodological criticism that their estimates are biased (with the potential exception of the intergenerational curve of the median). The purpose of this section is to now consider *other* methodological problems that arise under conditions (a) or (b).

The main additional problem with LS's proposed approach, under either scenario (a) or (b), is that it is inconsistent with the stricture that measures of intergenerational persistence should be insensitive to scale. The use of scale-invariant measures is widely endorsed not only in research on intergenerational mobility and persistence but also in research on economic inequality more generally. The widely accepted axiom of scale invariance, as it is labelled within the economic inequality field, posits that equi-proportional changes of incomes should not affect our inequality measures (e.g., Cowell 2000). The field of intergenerational mobility and persistence, which studies inequality associated with economic origins, has not referred to this property with the same name (i.e., scale invariance), but it has still been widely embraced as desirable and indeed all major approaches to measuring income mobility that purport to address inequality of opportunity are scale invariant. Most importantly, intergenerational elasticities have this property, as stressed by Mulligan (1997:24-25) among many others. This property is also satisfied by the measure of income-share mobility (Bratberg et al. 2017), the intergenerational linear correlation (e.g., Nybom and Stuhler 2017), measures of transition probabilities across income quintiles or deciles (e.g., Jäntti et al 2006), and the rank-rank slope and the expected rank at the 25^{th} percentile of parental income (e.g., Chetty et al. 2014; Corak, Lindquist, and Mazumder 2014).

By contrast, LS's proposed approach does not satisfy this property, as Figure A1 reveals. The left panel of Figure A1 reproduces LS's Figure 2, while the right panel of Figure A1 is based on the same data as the left panel except that children's income has now been multiplied by 1.6. The figure suggests that origins are much more consequential in the second population than in the first. When LS described the left panel, they concluded that the "line for the 90th percentile indicates that the probability of having a particularly high offspring income is quite sensitive to small changes in parent income" (p. 9). We might be tempted to conclude, on the basis of comparing the two panels, that there is yet more "sensitivity" to

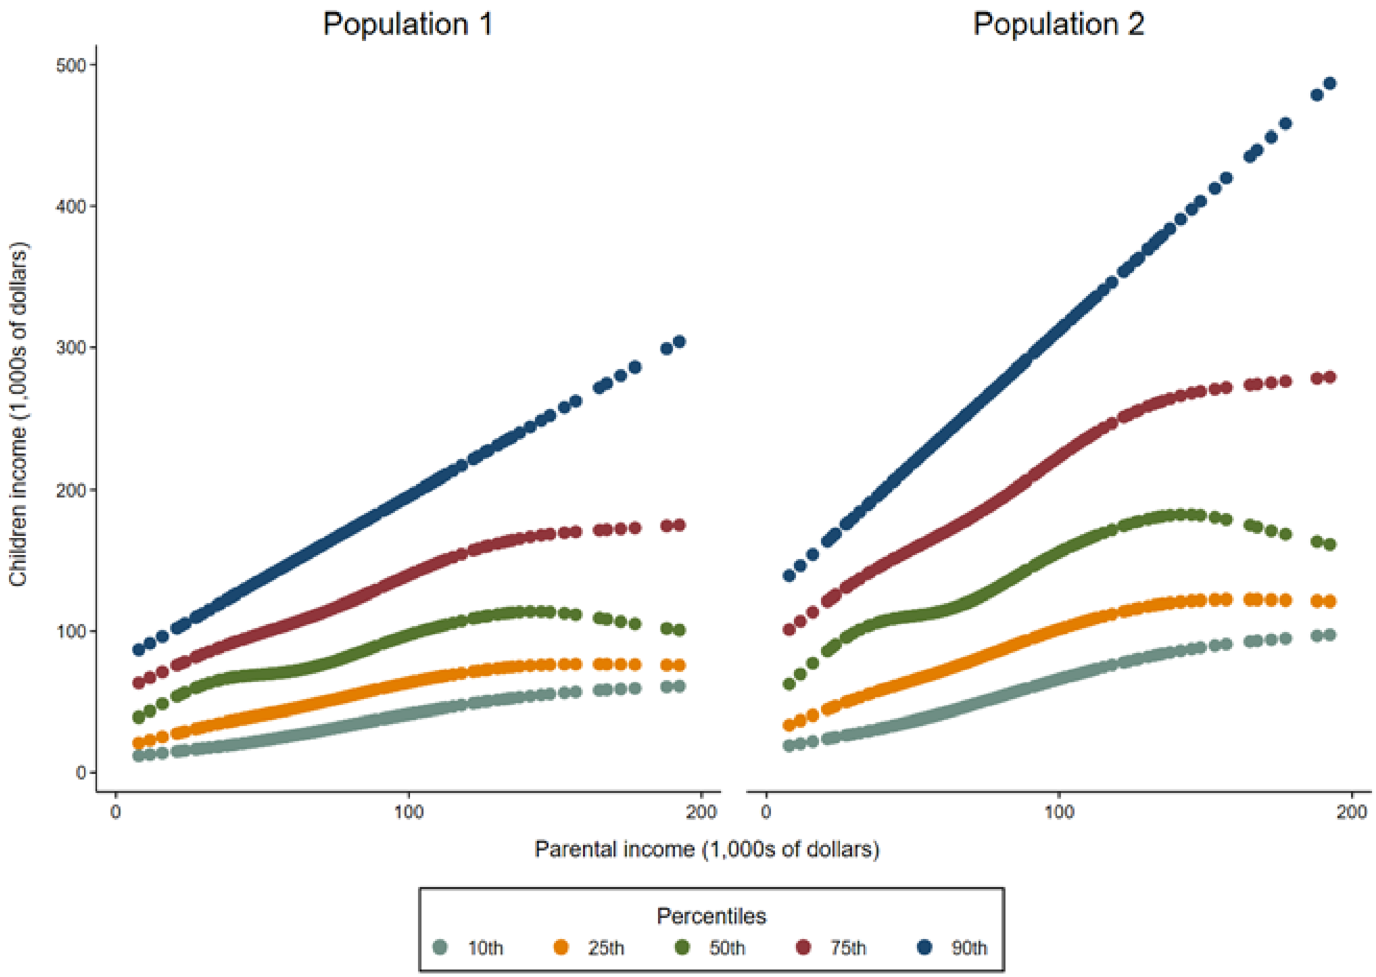

Figure A1. Quantile intergenerational curves of family income for two populations in the space spanned by the income variables. Although the figure suggests that people's economic origin is much more consequential in the second population than in the first, the inequality across conditional distributions at any quantile, and the share of income inequality transmitted from parents to children, is the same for both populations. The income data underlying the curves on the right panel are the same as the data underlying the curves on the left panel, but children's income has been multiplied by 1.6.

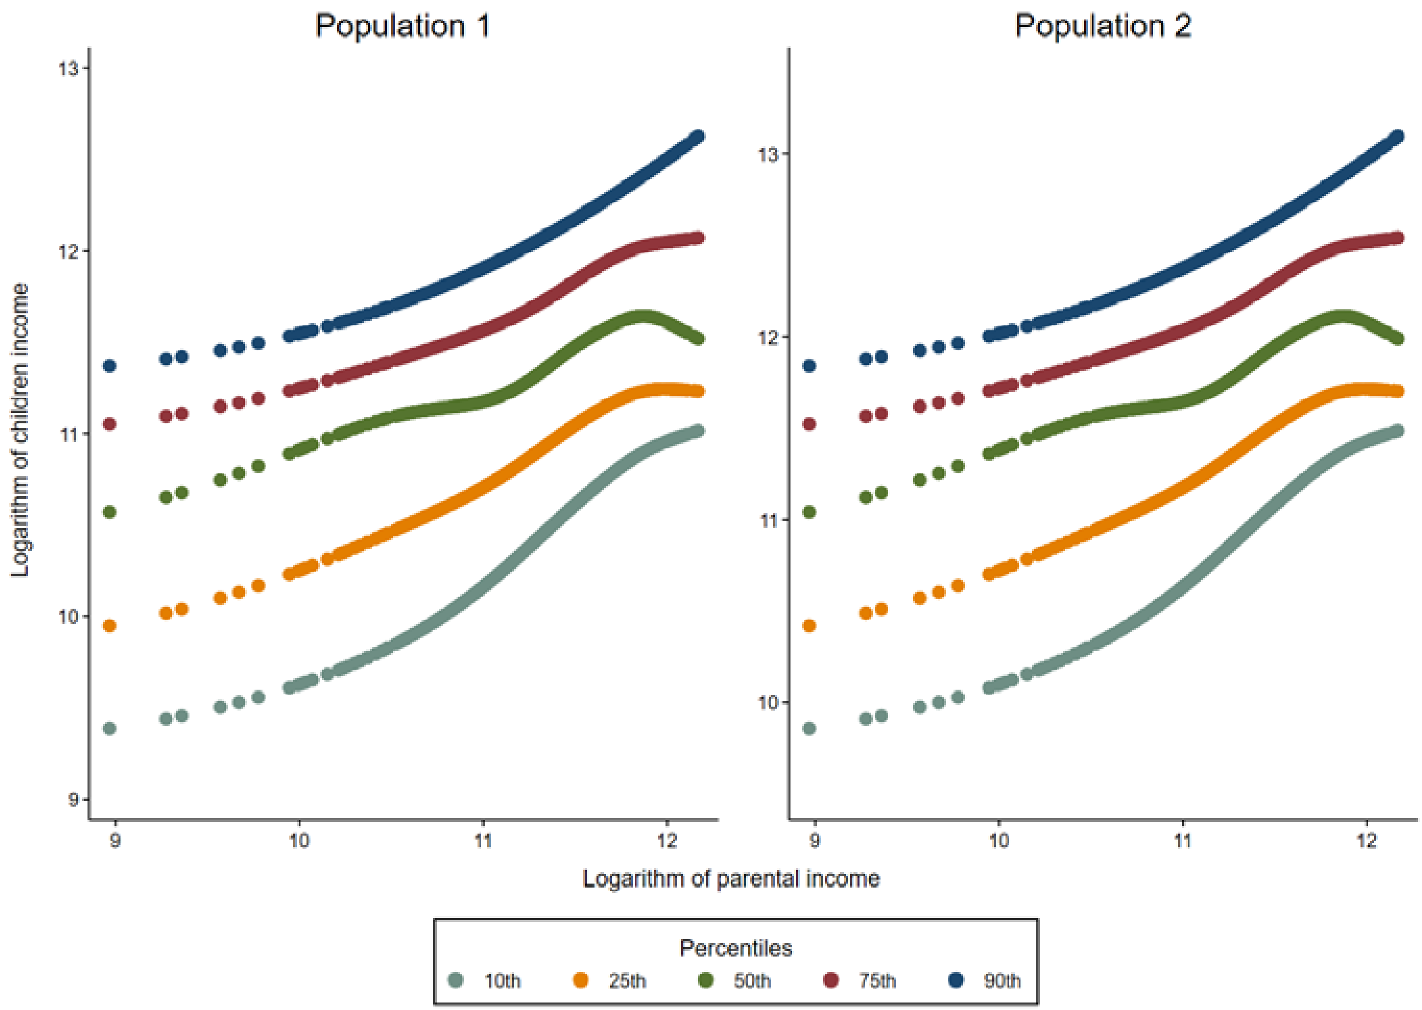

Figure A2. Quantile intergenerational curves of family income for two populations in log-log space. The data underlying the curves are the same as in Figure A1. The representation of the quantile curves in log-log space makes clear that economic origin is equally consequential in both populations.

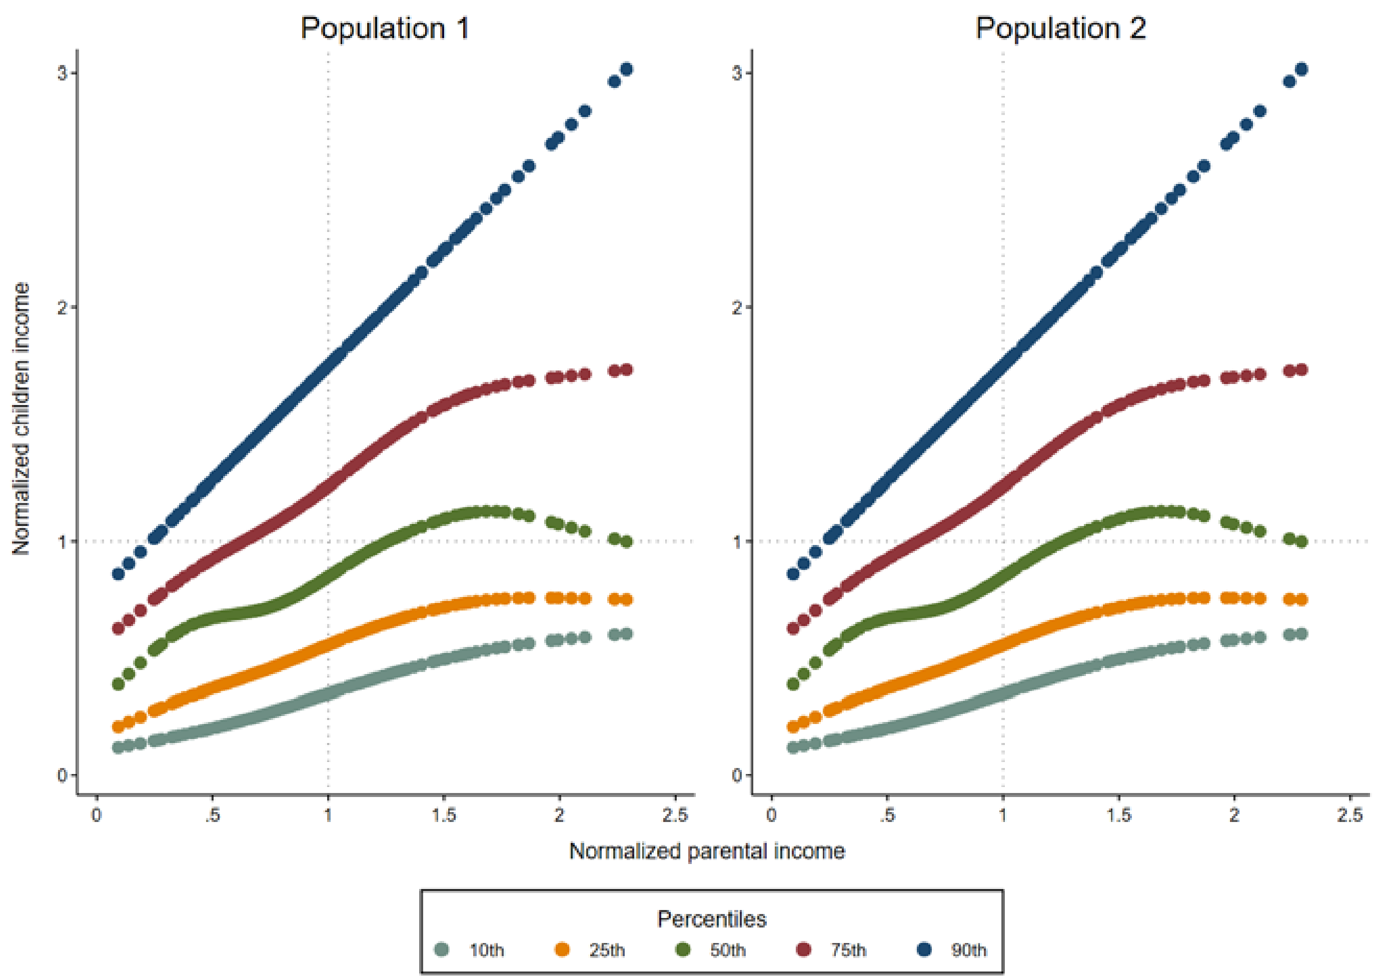

Figure A3. Quantile intergenerational curves of family income for two populations, with normalization. The data underlying the curves are the same as in Figure A1. The normalization involves dividing each income variable by its mean. The visual representation of normalized quantile curves also makes clear that economic origin is equally consequential in both populations.

parent income within the second population. However, because the only difference between the left and right panels is that, for the right panel, children's income has been multiplied by 1.6, we know that (a) the inequality across conditional distributions at any quantile is the same for both populations, and (b) the share of income inequality transmitted from parents to children is the same in both populations as well.[1]

This illustrates that LS's proposed approach does not satisfy scale invariance and that, as a result, it cannot be used to pursue the goals that research oriented toward inequality of opportunity has typically pursued. It bears noting, moreover, that these types of cross-population comparisons are not only relevant when the goal is to study temporal or cross-national variability in economic mobility. To the contrary, even researchers who focus squarely on one country or period have typically interpreted their findings by comparing them to what is known about other countries or periods, a comparison that becomes impossible with LS's proposed approach.

We conclude this section on a more positive note. It is possible to restore scale invariance by (a) representing quantile intergenerational curves in log-log space (see Mitnik et al. [2018] for a relevant example with mean rather than quantile intergenerational curves), (b) presenting quantile intergenerational curves of normalized income (where the normalization is achieved, for instance, by dividing the income of children and parents by the mean income of each generation, as in Bratberg et al. 2017), or (c) estimating quantile intergenerational curves of income ranks rather than income levels (similar to the expected-rank intergenerational curves presented in Chetty et al. 2014). Figure A2 shows the curves from Figure A1 in log-log space, while Figure A3 shows them after normalizing income (by dividing it by mean income). These types of representations will produce quantile intergenerational curves with identical shapes regardless of "income scale" (as long as the axes have identical lengths and the curves are centered similarly with respect to the axes).

Alternative single-value measures

In addition to visual representations of quantile intergenerational curves, LS also propose a new family of one-value measures of intergenerational persistence. It is useful to quote their proposal at length:

> A key selling point of [IGEs] . . . is that they offer single number summaries. These can be helpful for making comparisons across countries or time. If one wants such a summary and is comfortable with the ensuing information loss, it is possible to generate one by summarizing the fitted smooth curve. For instance, we can produce a first difference for (e.g.) the 50th percentile: (1) calculate the 50th percentile of offspring incomes at each observed parent income, (2) calculate the 50th percentile when $10k is added to each parent income, and then (3) report the population-average difference between (2) and (1). Median offspring income is (on average) $4k higher when parent income is $10k higher. Conducting the same first difference at the 90th

[1] These two conclusions assume the use of a measure of inequality that satisfies scale invariance. As noted earlier, all inequality measures typically used to study income inequality are scale invariant, including the Gini coefficient and its many variants, Theil's index and other entropy-based measures, the standard deviation of log income, and the relative mean deviation. We define the share of transmitted inequality as the ratio between inequality in conditional expected income and overall inequality in the parents' generation (see Mitnik et al. 2019). The same result obtains when, instead of using the conditional expectation, one uses other indices to assign values to conditional distributions (e.g., LS's preferred index, the conditional geometric mean).

> percentile produces a higher first difference of $12k, reflecting the steeper slope of the 90th percentile curve (p. 9).

With this passage, LS are now referring directly to the task of carrying out comparisons across countries or time periods, meaning that our comments from the previous section carry special force. To illustrate just how problematic their proposal is, let's consider a toy example in which (a) we are examining change within a country across a 30-year period, (b) we rely on the median intergenerational curve to characterize that change (given that it might be, as noted in the main text, a consistent estimate of the long-run curve even if based on short-run income), (c) all income values have been adjusted with a consumer price index that measures inflation perfectly over that period, (d) the children in the first period are the parents in the second period, (e) the shapes of the income distributions are the same in all cases (e.g., Dagum distributions with the same shape parameters), and (f) average real income is the same for parents and children in the first period but is 60 percent higher for children in the second period. This situation is equivalent to that shown in Figure A1, with the left panel representing the first period, and the right panel representing the second period. If we were to use LS's proposed one-value measure to conduct our comparison, we would wrongly conclude that the share of economic inequality transmitted across generations has increased 60 percent between the two periods (with the conditional median indexing the value of conditional distributions). But in fact we know that the share is the same in both cases (because we have imposed an equi-proportional change in income). More generally, any comparative results under LS's proposed one-value measure would be highly sensitive to the choice of price indices, purchasing power parities, and the arbitrary values used for the computation of first differences (e.g., computing those differences after adding $10,000, rather than $5,000 or $20,000, to parental income). We find it difficult to imagine a more unattractive proposal.

Is there a fix available? There indeed is. If the median intergenerational curve is represented in log-log space, the expected slope of this curve across values of parental income would make for a fine one-value measure. This is of course nothing other than a proposal to nonparametrically estimate the IGE-M. The expected arc elasticity of the conditional median across all possible pairs of values of parental income is another viable option (see Mitnik et al. [2018:26-28] for relevant work).

Nonnegative income variables and measurement error

At the very beginning of their comment, LS make the puzzling claim that "the zeros problem in empirical work arguably arises only due to measurement error since many (perhaps all) people have non-zero lifetime incomes" (p. 2). Are LS truly suggesting that it is a case of measurement error when a woman marries young, remains married until death, and thus has no lifetime earnings (because she never works outside her home)? Are they suggesting that we should value domestic labor at the market-price equivalent, impute "earnings" accordingly, and then pretend that "domestic workers" actually receive such pseudo-earnings? Are they suggesting that, because lifetime domestic workers will almost always have nonzero family income, we can solve the zeros problem by abandoning the study of earnings mobility and focusing exclusively on family-income mobility? Are they postulating an imaginary country in which a job is provided to anyone who wants one (and that, moreover, everyone "wants one")? It is hard to know how to choose among these interpretations because they are all implausible.

But it gets worse. In their follow-up to this statement, LS go on to note that "the present comment sets this concern aside and accepts MG's contention that the variables of interest truly have some zero values that are not simply measurement error" (p. 2). This statement betrays an abject misunderstanding of the zeros problem. Although the frequency of lifetime zeros is declining for women in the United States and other countries, many men and women will nonetheless have zeros on the short-run measures that are nearly always used for estimation. This is not a problem that will go away. The field will always have to rely on such short-run measures because there is a compelling interest in assessing lifetime earnings mobility without waiting for cohorts to age out of the labor force (and thus complete their lifetime earnings). It is very strange in this context to suggest that we are simply *contending* that there are zeros, that this is a contention that can be plausibly doubted, and that we should be grateful that LS are willing to set their doubts about this contention aside.

The next statement in this passage has LS encouraging "further work that considers the underlying measurement construct and its relationship to the measured outcome" (p. 2). It is difficult to evaluate this programmatic suggestion because LS make no reference to the large existing literature on this topic and the aspects of it that they believe need further development. By contrast, we engaged with this literature at length, dedicating a full section of our article (see Mitnik and Grusky 2020, Section 3) to the methodological problems that arise when a mobility measure defined in terms of long-run income has to be estimated with short-run proxy measures. The strategies that may be used to address some of those problems are further discussed in Section 6 and in Online Appendix B of our paper. In their own empirical analysis, LS ignored the issue of the relationship between short-run income measures and the long-run measures of interest (as we have discussed in the main text), a puzzling analytic decision because it is inconsistent with their own injunction to attend to measurement issues. Although "further work that considers the underlying measurement construct and its relationship to the measured outcome" (p. 2) is always welcome, any such work needs to start by examining the literature on generalized-error-in-variables and other measurement-error and related models, much of which we cite in our paper (e.g., Böhlmark and Lindquist 2006; Jerrim et al. 2016; Haider and Solon 2006; Mazumder 2001, 2005; Mitnik 2017, 2019, 2020; Nybom and Stuhler 2016; Solon 1992: Appendix).

B. Is conditional median income an attractive basis for a workhorse intergenerational elasticity?

In Mitnik and Grusky (2020), we call for replacing the IGE-G with the IGE-E, not with the IGE-M. We do so, in part, because we concluded that the IGE-M is poorly suited to play the role of the field's "workhorse intergenerational elasticity." The three reasons for this conclusion are briefly laid out here.

The IGE-M cannot be used to disentangle the different channels for transmitting economic status across generations

In Mitnik and Grusky (2020), we show that it is possible to use the IGE-E, but not the IGE-G, to study the channels (e.g., labor market, marriage market) through which the intergenerational transmission of advantage occurs. We introduce an expression (equation 19) showing how the family-income average elasticity across values of parental income depends on (1) the average elasticity of the expectation of the child's own income, (2) the average elasticity of the expectation of the spouse's income conditional on marriage, and (3) the average elasticity of the probability of marriage (where all elasticities are with respect to the child's parental income and all expectations are with respect to the distribution of the child's parental income). A comparable expression cannot be derived with the IGE-M because, at any given value of parental income, median family income cannot be expressed as a function of median own income and median spousal income, nor is it possible to express the median spousal income as a function of median spouse's income conditional on marriage and the probability of being married.

It has proven much harder to analyze the effects on estimates of estimating the IGE-M with short-run proxy income variables instead of the long-run income variables of interest

Generalized error-in-variables models have been developed and empirically tested for both the IGE-G and the IGE-E (Haider and Solon 2006; Nybom and Stuhler 2016; Böhlmark and Lindquist 2006; Mitnik 2017, 2019, 2020). These measurement-error models (a) address the effects of lifecycle and attenuation biases that result from the estimation of the IGE-G and IGE-E with short-run income variables and, when relevant, also the bias associated with the use of the invalid instruments typically available to mobility scholars, (b) suggest strategies to eliminate those biases when generating point estimates or, when that is not possible, to bound the IGEs or to combine estimators biased in opposite directions to produce set estimates, and (c) provide guidance in the interpretation of results. Developing similar measurement-error models to estimate the IGE-M has proved more difficult. In the absence of any formal model, it would still be possible in some data contexts to conduct empirical analyses aimed at improving our understanding of the role of measurement error, but so far such empirical analyses have not been carried out.

The IGE-M is more difficult to estimate with the statistical packages widely used by social scientists and, as long as the estimates rely on short-run income variables or invalid instruments, they are difficult to interpret

In Mitnik and Grusky (2020: Sec. 6.3), we distinguish several data contexts in which the IGE-G has been parametrically estimated, and we then show that there are estimators available to easily estimate the IGE-E in those contexts. The following points are relevant here: (a) the Poisson pseudo-maximum likelihood (PPML) estimator (Santos Silva and Tenreyro 2006) may be used to estimate the IGE-E in all situations in which the IGE-G has been estimated with the ordinary least squares (OLS) estimator; (b) the additive-error version of the generalized method of moments (GMM) instrumental-variable (IV) estimator of the Poisson or exponential regression model (Mullahy 1997; Windmeijer and Santos Silva 1997) may be used to estimate the IGE-E in all situations in which linear IV estimators (e.g., the two-stage least squares estimator) have been used to estimate the IGE-G; (c) the PPML and GMM-IVP estimators can be combined to set estimate the IGE-E in all situations in which the OLS and a linear IV estimator have been combined to set estimate the IGE-G; (d) a two-sample GMM estimator of the exponential regression model (Mitnik 2017), or GMM-E-TS estimator, can be used to estimate the IGE-E in all situations in which the two-sample two-stage least squares (TSTSLS) estimator has been used to estimate the IGE-G; (d) all the estimators of the IGE-E just mentioned may be used both with equal-probability samples and with samples that are the result of complex sampling designs (e.g., samples requiring the use of sampling weights in estimation); and (e) all the estimators of the IGE-E are available in at least one of the statistical packages broadly used by social scientists, while the PPML estimator (the "core estimator" aimed at playing the role that the OLS estimator has played for the IGE-G) is available in all major statistical packages used by social scientists.

For the IGE-M, a first goal is to secure an estimator able to play the core role that the OLS and PPML estimators play for the IGE-G and IGE-E. We thus want an estimator of the following model:

$$Med(Y|x) = \exp(\gamma_0 + \gamma_1 \ln x),$$

where X is parental income, Y is children's income, Med is the median operator, and γ_1 is the constant IGE-M (here we assume the elasticity is constant to simplify the discussion). Estimating this model is the best option, all else equal, because it is a valid approach even when the data include zeros, exactly as is the case with the IGE-E. Although the model can be estimated using the algorithm introduced by Koenker and Park (1996), to the best of our knowledge the only available "canned routine" for estimating it is found in R (where it can be estimated with the function nlrq of package quantreg). However, estimation with this function is only possible with equal-probability samples.

The alternative is to estimate the model:

$$Med(\ln Y \,|x) = \gamma_0 + \gamma_1 \ln x,$$

after replacing any 0s in Y by 1s (or by a positive value smaller than the minimum positive value of Y in the dataset, if the latter is not larger than 1). This approach is valid, as it relies on the equivariance-to-monotone-transformations property of quantiles, which entails that $Med(\ln Y \,|x) = \ln Med(Y|x)$. Although estimation of this model is a simple task with all major statistical packages used by social scientists, support for complex survey data is limited. For instance, while the model can be estimated using sampling weights with any major statistical package, only some of them can factor in the existence of clusters for computing confidence intervals (and in some cases only by relying on computer-intensive resampling methods). And, of course, we do not have anything close to a good understanding of the effects on estimates of using short-run proxy income variables to estimate the model. We also know little about the strategies that may be used to reduce the biases that relying on short-run measures can be expected to generate.

It is not clear which estimator for the IGE-M can successfully play the role that the IV linear estimators and the GMM-IVP estimator play for the IGE-G and IGE-E, respectively. A potential candidate is an instrumental-variable estimator developed by Chernozhukov et al. (2015), which has been implemented in one major statistical package (Stata), can be used with unequal-probability samples and allows to compute cluster-adjusted robust standard errors. The assumption that the estimator produces upward-biased estimates with the invalid instruments typically available to mobility scholars (e.g., parental education) would be based, for the time being at least, on a heuristic argument, as neither a formal analysis of the matter nor any empirical evidence is available.

Lastly, we are only aware of one two-sample estimator that could be used to estimate the IGE-M in the data context in which the TSTSLS and GMM-E-TS estimators can be used to estimate the IGE-G and the IGE-E, respectively. This two-sample estimator (see Grawe 2004) relies on strong distributional assumptions and, to the best of our knowledge, is not available in any of the statistical packages broadly used by social scientists. For these two reasons, it is not an attractive estimator. It is also unfortunate that no formal analysis or empirical evidence exists regarding the sign of the bias of the estimator when used with the invalid instruments typically available to mobility scholars.

References

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.

Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel Schnitzlein and Kjell Vaage. 2017. "A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the U.S." *The Scandinavian Journal of Economics* 119 (1): 72–101.

Chernozhukov, Victor, Iván Fernández-Val, and Amanda Kowalski. 2015. "Quantile Regression with Censoring and Endogeneity." *Journal of Econometrics* 186 (1): 201–221.

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.

Corak, Miles, Matthew Lindquist and Bhashkar Mazumder. 2014. "A Comparison of Upward and Downward Intergenerational Mobility in Canada, Sweden and the US." *Labour Economics* 30: 185–200.

Cowell, 2000. "Measurement of Ineuality." In *Handbook of Income Distribution*, Vol. 1, edited by A. B. Atkinson and F. Bourguignon. New York: Elsevier.

Grawe, Nathan. 2004. "Intergenerational Mobility for Whom? The Experience of High- and Low-Earnings Sons in International Perspective." *Generational Income Mobility in North America and Europe*, edited by Miles Corak. Cambridge: Cambridge University Press.

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

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, Álvaro Choi, Rosa Simancas. 2016. "Two-Sample Two-Stage Least Squares (TSTSLS) Estimates of Earnings Mobility: How Consistent are They?" *Survey Research* Methods 10 (2): 85-102.

Koenker, Roger and Beum Park. 1996. "An Interior Point Algorithm for Nonlinear Quantile Regression." *Journal of Econometrics* 71 (1-2): 265-283.

Lundberg, Ian and Brandon Stewart. 2020. "Comment: Summarizing Income Mobility with Multiple Smooth Quantiles Instead of Parameterized Means." *Sociological Methodology* 50 (1): ###-###.

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.

Mitnik, Pablo. 2017. "Two-Sample Estimation of the Intergenerational Elasticity of Expected Income." Stanford Center on Poverty and Inequality Working Paper.

Mitnik, Pablo. 2019. "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. 2020. “Intergenerational Income Elasticities, Instrumental Variable Estimation, and Bracketing Strategies.” *Sociological Methodology* 50 (1): ##-##.
Mitnik, Pablo, Victoria Bryant, 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, Victoria Bryant, and Michael Weber. 2019. “The Intergenerational Transmission of Family-Income Advantages in the United States.” *Sociological Science* 6(15): 380-415.
Mitnik, Pablo and David Grusky. 2020. “The Intergenerational Elasticity of What? The Case for Redefining the Workhorse Measure of Economic Mobility.” *Sociological Methodology* 50 (1): ##-##.
Mulligan, Casey. 1997. *Parental Priorities and Economic Inequality*. Chicago: The University of Chicago Press.
Mullahy, John. 1997. “Instrumental-Variable Estimation of Count Data Models: Applications to Models of Cigarette Smoking Behavior.” *Review of Economics and Statistics* 79(4): 586-593.
Nybom, Martin and Jan Stuhler. 2016. “Heterogeneous Income Profiles and Life-Cycle Bias in Intergenerational Mobility Estimation.” *The Journal of Human Resources* 51(1): 239-268.
Nybom, Martin and Jan Stuhler. 2017. “Biases in Standard Measures of Intergenerational Income Dependence.” *The Journal of Human Resources* 52 (3): 800-825.
Santos Silva, J. M. C. and S. 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: 393-408.
Windmeijer, F. A. G. and J. M. C. Santos Silva. 1997. “Endogeneity in Count Data Models: An Application to Demand for Health Care.” *Journal of Applied Econometrics* 12: 281-294.