"PERSON-WEIGHTED" VERSUS "DOLLAR-WEIGHTED":
A FLAWED CHARACTERIZATION OF TWO
INTERGENERATIONAL INCOME ELASTICITIES

Pablo A. Mitnik
(pmitnik@stanford.edu)

September, 2017

Acknowledgements

Abstract

The intergenerational income elasticity (IGE) has been widely misinterpreted as pertaining to the conditional expectation of children's income when it pertains to its conditional geometric mean, and is affected by serious methodological problems. This has led to a call to replace the conventionally estimated IGE by the IGE of the expectation. This paper shows that a contrasting characterization of these two IGEs as, respectively, 'person-weighted' and 'dollar-weighted' elasticities, which casts doubts on the desirability of the proposed replacement, is the joint result of a category mistake—equating quantile-specific elasticities to person-specific elasticities—and of misconstruing the nature of the conventional IGE and the epistemic goal it has been meant to serve.

Introduction

  The intergenerational elasticity (IGE) has long been the workhorse measure of economic mobility (e.g., Solon 1999; Black and Devereux 2011). Nevertheless, Mitnik and Grusky (2017) have recently shown that the IGE has been widely misinterpreted. While it is assumed to refer to the expectation of the children's earnings or income conditional on their parents' income, it pertains in fact to the conditional geometric mean of the children's earnings or income; this invalidates all the interpretations that have typically been imposed on its estimates, including the archetypical interpretation as a measure of regression to the arithmetic mean.

  If this were the only problem affecting the conventionally-estimated IGE, mobility scholars could perhaps address it by changing the way they interpret their IGE estimates. Unfortunately, the unwitting reliance on the geometric mean also generates very serious methodological problems, which do not have any easy solution (Mitnik and Grusky 2017). Both the conceptual and the methodological problems, however, can be solved at once by replacing the IGE of the geometric mean (the de facto estimated IGE) by the IGE of the expectation (the IGE that mobility scholars thought they were estimating) as the workhorse intergenerational elasticity. Mitnik and Grusky (2017) have called for effectuating such replacement.

  Although the case for such replacement is seemingly very strong, Chetty, Hendren, Kline and Saez (2014) have advanced a contrasting characterization of the two elasticities—the conventional IGE and the IGE of the expectation—that casts doubts on its desirability. They argued that while the former IGE is a "person-weighted" IGE that weights all individuals equally, the IGE of the expectation is a "dollar-weighted" IGE that weights individuals in proportion to their income. This suggests that the IGE of the expectation is conceptually inadequate given the purposes pursued by mobility scholars when using the conventional IGE.

The main goal of this paper is to show that the contrasting characterization of the two IGEs advanced by Chetty et al. (2014) is the joint result of a category mistake—equating quantile-specific elasticities to person-specific elasticities—and of misconstruing the nature of the conventional IGE and the epistemic goal it has been meant to serve. The rest of the paper has four sections. I first summarize Mitnik and Grusky's (2017) arguments for redefining the IGE. Next, I present Chetty et al.'s comparative characterization of the two IGEs and my criticism of it. The last section offers some concluding remarks.

The IGE of what? Redefining the workhorse intergenerational elasticity

The standard population regression function (PRF) posited in the literature, which assumes the elasticity is constant across levels of parental income, is:

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

where Y is the child's long-run income or earnings, X is long-run parental income or father's earnings, β_1 is the IGE as specified in the literature, and I use expressions like "Z|w" as a shorthand for "$Z|W = w$." The parameter β_1 is not, in the general case, the elasticity of the conditional expectation of the child's income, as mobility scholars have widely assumed. This would hold as a general result only if $E(\ln Y|x) = \ln E(Y|x)$. But, due to Jensen's inequality, the latter is not the case. Instead, as $E(\ln Y \mid x) = \ln \exp E(\ln Y \mid x)$, and $GM(Y|x) = \exp E(\ln Y \mid x)$, Equation [1] is equivalent to

$$\ln GM(Y|x) = \beta_0 + \beta_1 \ln x, \qquad [2]$$

where GM denotes the geometric mean operator. Therefore, β_1 is the elasticity of the conditional geometric mean, i.e., 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.[1]

---

[1] The parameter β_1 is (also) the IGE of the expectation only when the error term satisfies very

As the geometric mean is undefined whenever an income distribution includes zero in its support, the IGE is undefined as well when this is the case. Mitnik and Grusky (2017) have shown that this has serious methodological consequences—in a nutshell, it badly hinders the study of gender and marriage dynamics in intergenerational processes, and leads to IGE estimates affected by substantial selection biases.

To address these problems, Mitnik and Grusky's (2017) have called for redefining the workhorse measure of economic mobility. This entails replacing the PRF of Equation [1] by:

$$\ln E(Y|x) = \alpha_0 + \alpha_1 \ln x, \qquad\qquad [3]$$

where $Y \geq 0, X > 0$ and $\alpha_1 = \frac{d \ln E(Y|x)}{d \ln x}$ 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. Crucially, (a) all interpretations incorrectly applied to the conventional IGE are correct under this formulation (Mitnik and Grusky 2017:Section 5.A), and (b) the IGE of the expectation is fully immune to the methodological problems affecting the IGE of the geometric mean and, in particular, is very well suited for studying the role of marriage in the intergenerational transmission of advantage (Mitnik and Grusky 2017:Section 5.B).

The IGEs as 'person-weighted' and 'income-weighted' elasticities

The distinction between person-weighted and income-weighted elasticities is not new. It has played an important role in the field of public finance, where it has been shown that the policy-relevant elasticity of taxable income with respect to tax rates is the income-weighted elasticity (see, e.g., Saez, Slemrod and Giertz 2012). The following example illustrates well the difference between the two types of elasticities and their role in public finance:

---

special conditions (Santos Silva and Tenreyro 2006; Petersen 2017; Wooldridge 2002:17).

[A]n income-weighted ETI [elasticity of taxable income] should reflect the percent change in total taxable income (associated with a 1% increase in the NTR [net-of-tax rate])—instead of the average of individual percent changes. For example, consider two taxpayers: Person 1 has income of $10,000, and person 2 has income of $1 million. In response to a 1% decrease in the NTR, suppose that person 1 reduces his income (by 0.2%) to $9,980 and that person 2's income falls (by 1%) to $990,000. In this instance, the unweighted ETI [i.e., the person-weighted measure] equals 0.60 …, whereas the income-weighed measure is nearly two-thirds larger, at 0.99 . . . (Giertz 2010:419).[2]

Chetty et al.'s (2014) comparative characterization of the conventional IGE and the IGE of the expectation posits that the distinction between person-weighted and income-weighed elasticities is also relevant in the intergenerational mobility field. Let $Q_Y(x, \tau)$ denote the τ^{th} quantile of the conditional distribution of the child's long-run income when long-run parental income equals x. Chetty et al. (2014) showed that, at any value of parental income x, the conventional IGE and the IGE of the expectation can be written as follows:

$$\frac{dE(\ln Y \,|x)}{d \ln x} = \int_0^1 \frac{d \ln Q_Y(x, \tau)}{d \ln x} \, d\tau \qquad [4]$$

$$\frac{d \ln E(Y|x)}{d \ln x} = \int_0^1 \frac{Q_Y(x, \tau)}{E(Y|x)} \frac{d \ln Q_Y(x, \tau)}{d \ln x} \, d\tau. \qquad [5]$$

From the representations in Equations [4] and [5], Chetty et al. (2014) argued that the conventional IGE "can be interpreted as the average elasticity of child income with respect to parent income in a model with heterogeneous elasticities, while Mitnik et al.'s [2015] new measure [the IGE of the expectation] is a dollar-weighted (i.e., child-income-weighted) average of the same elasticities" [one typo corrected]. And, because of this, they also claimed that the conventional IGE "weights all individuals with positive income equally" and referred to it as a "person-weighted" IGE, while they referred to the IGE of the expectation as a "dollar-weighted" elasticity where the weights "are an increasing function of the child's income" (Chetty et al.

---

[2] Here, $0.60 = \left[\frac{1\% + 0.2\%}{2}\right] / 1\%$, while $0.99 = \left\{100 \left[\frac{20 + 10,000}{10,000 + 1,000,000}\right] \%\right\} / 1\%$.

2014:1574 and Online Appendix C).[3]

The foregoing entails that the IGE on which mobility scholars have been interested is an average of person-level IGEs while the IGE of the expectation is something else, which in turn suggests that (a) the IGE of the expectation is conceptually inadequate given the purposes pursued by mobility scholars when using the conventional IGE, and (b) we may be better served by redoubling our efforts to solve the methodological problems affecting the conventional IGE than by replacing it by the IGE of the expectation, as Mitnik and Grusky (2017) have proposed.

Criticism of a flawed characterization

Let's start with what Chetty et al.'s (2014) characterization implies regarding the nature of the conventional IGE and the epistemic goal it is supposed to serve. As a person's long-run parental income is a fixed attribute of that person, a person-specific intergenerational elasticity is by necessity a causal parameter, that is, it cannot simply reflect an empirical association across people, or across times "within" a particular individual.[4] Therefore, Chetty et al.'s characterization entails that the conventional IGE is also a causal parameter: An average of marginal proportional effects resulting from marginal proportional treatments.

This is clearly inconsistent with the way in which the conventional IGE has been understood in the mobility literature. Indeed, because of the correlation between parental income

---

[3] Section C of Chetty et al.'s (2014) Online Appendix was written as a response to an early draft of Mitnik et al. (2015). That draft did not characterize the conventional estimand as the IGE of the geometric mean; it just pointed out it was not the IGE of the expectation in the general case.

[4] This is, of course, consistent with how the notion of person-based elasticity has been used in public finance to refer to a behavioral elasticity that measures how a person's taxable income changes when he or she responds to a change in the tax rate by changing his or her behavior, e.g., his or her hours of work.

and a large number of causally relevant factors not included in specifications like that of Equation [1]—which can be thought of as omitted variables that have been very purposefully omitted—mobility scholars have not interpreted the IGE as a parameter measuring the causal effect of parental income. Rather, they have conceived of it as a descriptive measure—comparable in nature to, e.g., the Gini coefficient—providing information on how much better the children's economic outcomes become as their parental economic status increases, and which therefore can be interpreted as an imperfect index of the advantages conferred by the circumstances of birth. For instance: (a) mobility scholars often refer to the IGE as a "descriptive statistic" (e.g., Aronson and Mazumder 2008:146; Hertz 2007:26), or to Equation [1] as capturing "a simple statistical relationship" (e.g., Stuhler 2012:2), (b) Bjorklund and Jantii (2011:Secs. 4 and 6) describe the IGE as providing information on the "association" between parental and children's income and point out that "the causal effect of parental income is conceptually different and much more difficult to estimate," and (c) the key premise of Black and Devereux's (2011) recent literature review in the *Handbook of Labor Economics* is that the IGE does *not* measure the causal effect of parental income or earnings. All this is, of course, compatible with the IGE being a function of various structural parameters (see, e.g. Solon 2004).

It follows that Chetty et al.'s (2014) characterization of the conventional IGE, which would make it a causal parameter, (a) is badly aligned with the epistemic goals that that IGE has been meant to serve, and (b) implies that each quantile-specific elasticity can itself be legitimately interpreted as a causal parameter, in spite of the fact that many variables known to be causally relevant are omitted from the analysis.

The latter fact suggests that something more fundamental may be amiss in Chetty et al.'s analysis, and that is indeed the case. Equation [4] does show that the conventional IGE is

equivalent to a simple average of all quantile-specific IGEs. But quantiles are not people or, more precisely, quantile-specific IGEs are not person-specific IGEs. As an individual would not necessarily remain—and most likely would not remain—in the same quantile of the conditional income distribution if his or her parental income were different, the proposed interpretation of Equation [4] is simply incorrect. In fact, it follows immediately from results by Sasaki (2015; see especially Lemma 1) on what is identified by the quantile partial derivative, that $\frac{d \ln Q_Y(x,\tau)}{d \ln x}$ at each quantile τ is a remarkably complex function of person-level elasticities (in which the person-level elasticities are very far from being weighted equally). The characterization of the conventional IGE as a *person*-weighted IGE, or as an IGE that weights *individuals* equally, is therefore an incorrect characterization.

It may still be concerning that the IGE of the expectation is equal to a weighted average of all the quantile-specific elasticities, with the weights increasing with the quantile. Does this mean that there is something to the notion that the IGE of the expectation is a "dollar-weighted" IGE while the conventional IGE is not? Not at all. Equations [4] and [5] can be rewritten as follows:

$$\frac{dE(\ln Y \mid x)}{d \ln x} = \int_0^1 \frac{x}{Q_Y(x,\tau)} \frac{dQ_Y(x,\tau)}{dx} \, d\tau. \qquad [6]$$

$$\frac{d \ln E(Y|x)}{d \ln x} = \frac{x}{E(Y|x)} \int_0^1 \frac{d \, Q_Y(x,\tau)}{dx} \, d\tau. \qquad [7]$$

In the representation provided by Equation [7], the IGE of the expectation is equal to a simple average of all quantile-specific derivatives, multiplied by a scaling or standardizing factor. By comparing Equations [7] and [8] one could therefore conclude—using Chetty et al.'s logic—that the conventional IGE is a dollar-weighted average of derivatives in which the weights are a decreasing function of children's income, while the IGE of the expectation is an average of those

same derivatives in which all persons receive the same weight.[5] But this contrasting characterization would be as uninformative as that based on Equations [4] and [5]. The reason is that the relevant "elementary units" for the descriptive goal at hand are neither quantile-specific derivatives nor quantile-specific elasticities—incorrectly assimilated by Chetty et al. to person-specific elasticities—but full conditional distributions of children's economic status, as measured by their income or earnings. Both IGEs are predicated on summarizing the information in those conditional distributions by using a measure of central tendency, i.e., the geometric mean or the expectation. And both measures of central tendency are means of measures of economic status "denominated in dollars," which always increase with changes in density in favor of larger incomes or earnings (something that is not true, for instance, of the median). Therefore, in a rather trivial sense, both IGEs are based on "dollar-weighted measures of central tendency." There is no interesting sense, however, in which one IGE is dollar-weighted while the other is not.

Concluding remarks

This paper has shown that the distinction between person-weighted and dollar-weighted elasticities, which is certainly important in a context in which the elementary units are individual-level "behavioral elasticities" (i.e., public finance) (a) does not map into the differences between the conventional IGE and the IGE of the expectation, and (b) more broadly, is not relevant in the intergenerational mobility context, in which elasticities are not meant to measure causal effects. It follows that Chetty et al.'s (2014) characterization of the two IGEs is invalid and, therefore, that Mitnik and Grusky's (2017) case for replacing the conventional IGE by the IGE of the expectation remains warranted.

---

[5] This point is based on remarks made by Joao Santos Silva, in personal communication.

Cited references

Aaronson, David and Bhashkar Mazumder. 2008. "Intergenerational Economic Mobility in the US: 1940 to 2000. *Journal of Human Resources* 43(1): 139-172.

Björklund, Anders and Markus Jäntti. 2011. "Intergenerational Income Mobility and the Role of Family Background". *The Oxford Handbook of Economic Inequality*, edited by B. Nolan, W. Salverda and T. Smeeding. Oxford: Oxford University Press.

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.

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.

Giertz, Seth. 2010. "The Elasticity of Taxable Income during the 1990s: New Estimates and Sensitivity Analyses." *Southern Economic Journal* 77(2), 406–433.

Hertz, Tom. 2007. "Trends in the Intergenerational Elasticity of Family Income in the United States." *Industrial Relations* 46(1): 22-50.

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 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.

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

Saez, Emmanuel, Slemrod, Joel, and Seth Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature* 50(1):3-50.

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

Sasaki, Yuya. 2015. "What Do Quantile Regressions Identify for General Structural Functions?" *Econometric Theory* 31:1102–1116.

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 in Europa*, edited by Miles Corak. Cambridge: Cambridge University Press.

Stuhler, Jan. 2012. "Mobility across Multiple Generations: The Iterated Regression Fallacy." Discussion Paper 7072. Bonn: Institute for the Study of Labor.

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