

## THE CANADIAN GEOGRAPHY OF INTERGENERATIONAL INCOME MOBILITY\*

*Miles Corak*

Intergenerational income mobility varies significantly across Canada, with the 266 Census Divisions in the 1986 Census clustering into five non-contiguous regions. Nine complementary indicators are calculated for each Census Division using administrative data on a cohort of men and women born between 1963 and 1970. Collectively these indicators underscore the importance of simultaneously examining different dimensions of intergenerational mobility and also show that higher mobility is most strongly associated with less income inequality in the bottom half of the income distribution.

‘Equality of opportunity’—how likely it is, how it varies and what causes it—is an important issue in many countries. In part this is because of significant increases in income inequality and the now generally accepted view that higher inequality is associated with less social mobility. Former US President Obama and his senior advisers have made reference to this by referring to the Great Gatsby Curve (Krueger, 2012; White House, 2013; Furman, 2016), as have the International Monetary Fund, the Organization for Economic Cooperation and Development and Central Bankers like Janet Yellen (2014), formerly of the U.S. Federal Reserve, and Mark Carney (2014) of the Bank of England. This attention warrants a certain caution. ‘Equality of Opportunity’ is a multifaceted concept that cannot be simply defined and measured and while income mobility—the dimension most closely related to many policy discussions—may be an important aspect, it clearly is not the whole story.

\* Corresponding author: Miles Corak, Department of Economics and Stone Center on Socio-Economic Inequality, The Graduate Center, City University of New York, 365 Fifth Avenue, New York, NY 10016, USA. Twitter @MilesCorak, Website MilesCorak.com, Email: [mcorak@gc.cuny.edu](mailto:mcorak@gc.cuny.edu)

*This paper was received on 18 July 2017 and accepted on 11 September 2018. The Editor was Kjell Salvanes.*

The first draft of this paper was written while I was a visiting professor of economics at Harvard University during the 2015/16 academic year on leave from the Graduate School of Public and International Affairs at the University of Ottawa in Canada. I am grateful to Yuri Ostrovsky for help in updating the data and producing the statistics upon which my analysis is based. This was possible with the cooperation of Statistics Canada through a cost recovery project and the support of the Social Sciences and Humanities Research Council of Canada through Insight grant number 230936-191999-2001. Mary-Ellen Maybee and Brigit Levac of Statistics Canada kindly made Shapefiles associated with the 1986 Census available, allowing me to map the data, which was done using R. I also thank Raj Chetty, Nathan Hendren, Branko Milanovic, Martin Nybom, Jan Stuhler, the editor and three anonymous referees for feedback and advice. The paper was originally circulated as University of Chicago, Human Capital and Economic Opportunity Working Paper 2017-043 and versions have been presented at the April 2016 workshop on Inequalities and Families, organised by the Research Group on Human Capital at the Université du Québec à Montréal; in May 2016 to the 41st Congress of the Association des Économistes Québécois held in Québec City, to the 15th Annual North American Basic Income Guarantee Congress held in Winnipeg and at the Northwestern University Global Inequality Workshop in Evanston, Illinois; in October 2016 to the Sixth Annual NYU/UCLA Tax Policy Symposium on Tax Policy and Upward Mobility held at UCLA, Los Angeles, California; in November 2016 as the Shoyama Lecture presented in Ottawa to the Department of Finance, Government of Canada; in February 2017 to the Canada Seminar at Harvard University; in March 2017 to the Workshop on Interdisciplinary Approaches to Inequality and Mobility held at Duke University; in June 2017 to the 51st Annual Meetings of the Canadian Economics Association held at St Francis Xavier University in Antigonish, Nova Scotia; as a keynote address to the Canadian Research Data Centre Network 2017 National Conference, held in November 2017 at McGill University; and in seminars at the Graduate Center City University of New York, the University of Ottawa, the University of Toronto, and the University of Waterloo. The data and supplementary information, as well as an Online Appendix examining the robustness of the findings, are available at <https://milesacorak.com/equality-of-opportunity/>.

I offer an analysis of only this dimension, the degree to which adult incomes are related to the incomes of the families in which children were raised; in other words, the extent to which economic position in one generation echoes into the next. Even with this focus, there are a host of different statistics used to describe the process. My main contribution to the understanding of intergenerational income dynamics in Canada is to offer estimates of mobility across space in a way that pays attention to a broad suite of statistics capturing the many different ways academic researchers, policymakers and the general public view the process.

My primary objective is to estimate the degree of intergenerational income mobility within each of hundreds of subnational regions. The income tax forms Canadians submit to their governments is the information source, and offers very large sample sizes, in principle the entire population upon which I focus, those born between 1963 and 1970. Up to now the advantages of these data have been exploited to very precisely measure the degree of intergenerational mobility, to examine the process by income source and to estimate how it varies across the income distribution. All of this is for the country as a whole, yet the large sample holds the potential for more detailed regional analysis. I use these intergenerationally linked tax forms to calculate and examine different dimensions of economic mobility for each of the 266 Census Divisions defined in the 1986 Canadian Census.

My purpose is descriptive, to paint a comprehensive and detailed picture of the extent to which adult outcomes are related to family background and how the relationship varies across the country. Many of the factors and policies that theory suggests influence generational mobility vary significantly across space, or fall within the domain of provincial and municipal governments. Schooling, healthcare, employment opportunities for youth and basic income support for families involve not just the federal government, but in a major way provincial legislatures and city councils. A subnational portrait allows policymakers to know their own regions and to make comparisons with others. While the existing literature suggests that Canada is an intergenerationally mobile country, and in particular more mobile than the USA, a national focus potentially masks significant geographic variation within the country and places a blind spot over the experiences and needs of certain demographic groups and communities.

Intergenerational income mobility cannot be summarised with one statistic, yet the existing economics literature emphasises the intergenerational elasticity, the percentage difference in a child's adult income for a one percentage point difference in parent income. This is a valuable relative measure of mobility, offering a summary indicator of how income inequality evolves across generations. But it is not a complete description of the process, nor does it relate to some policy-relevant ways of thinking about it. Other measures are needed to speak to obvious concerns: to what extent does the current generation earn on average more or less than the previous generation, to what extent are rank and relative position transmitted across generations, and to what extent are low and top incomes transmitted intergenerationally? Income mobility, rank mobility and directional mobility—particularly upward mobility from the bottom—all relate to aspects of social welfare and capture the political imagination.

In addition to using a suite of statistics that collectively offer a broad picture of mobility, I also fully account for the role and outcomes of women. If only men work in the market, then family income is the income of the father. If the participation rates of women are low and intermittent as a result of childcare responsibilities falling on them, then their income is more easily approximated by the income of their partners. Focusing on fathers and adult sons, as much of the Canadian literature has done, makes the analysis easier by avoiding the need

to model these other dimensions and the impact of the marriage market. But these assumptions skew our understanding of the intergenerational transmission process and they potentially put aside significant numbers of individuals, both men and women, raised by single mothers. My analysis is based on 'family' income, loosely using that term to mean the total income of both partners, including periods when there is only one parent present. The analysis is conducted for sons and daughters without distinction, and as such reflects the host of influences on mobility, from labour market dynamics, to changes in participation rates, family formation and fertility.

I find that intergenerational income mobility varies across the country with a significant fraction of children raised by low-income parents facing considerable chances of an intergenerational cycle of poverty and limited opportunity of rising to the top. At the broadest level the Canadian landscape of economic opportunity should be thought of as being divided into five broad regions. These areas are not geographically contiguous, and provincial boundaries are not the dividing lines. My findings for Canada nuance the international comparative literature by suggesting that inequality in the bottom half of the income distribution is the dimension of inequality most relevant for understanding social mobility, correlating negatively with a comprehensive set of mobility outcomes.

## 1. Lessons from the Literature

I take three lessons from the theoretical and empirical literature that help frame my motivating questions and the development of an appropriate analytical data file. First, theory suggests that there are multiple determinants of intergenerational mobility, and therefore subnational differences in outcomes should be expected when these causal forces vary strongly across regions within a country. Relatedly, economic theory puts the focus on the intergenerational elasticity between parent and child incomes, but as valuable as this statistic is, it offers an incomplete guide on how to measure mobility. Second, recent research with American data stresses that documenting within-country variation may be particularly relevant for public policy, and by implication empirical researchers in other countries should seek to develop data appropriate for its study. Third, the Canadian literature has not examined within-country mobility differences even though existing research demonstrates that the available data seem well suited to the task.

### 1.1. *Theory and Statistical Indicators*

Becker and Tomes (1979; 1986) offer the workhorse model that motivates the analysis of the degree of 'regression to the mean' in incomes across generations, most commonly measured as the least squares estimate of the percentage change in a child's adult income associated with a percentage point change in the parent's income. This intergenerational income elasticity is a relative measure of intergenerational mobility, offering a summary indicator of the degree to which children tend to occupy the same position in the income distribution as their parents a generation earlier.

The Becker-Tomes model formalises the influence of 'inherited' characteristics, family investments in the human capital of their children, and the labour market payoff of these skills and characteristics on the earnings outcomes of children. This model has been refined in a number of

ways that can be interpreted as suggesting intergenerational mobility may vary not just over time, but also across space. As adapted by Becker *et al.* (2018) and Solon (2004; 2015) the model predicts that the intergenerational earnings elasticity will be higher when labour market inequality is higher. This is because more unequal labour markets—the differences in incomes reflecting rising returns to human capital—imply higher-income parents have both more resources and greater incentive to invest in the earnings capacity of their children, and to engage in other activities that give them a leg up in school and in finding jobs promoting their careers. This prediction applies equally to differences between communities across space, as it does to changes within a community through time.

The intergenerational elasticity may vary across space for other reasons. Communities may differ in the strength of their families—with respect to both monetary and non-monetary resources—and also in the structure of their labour markets and opportunities for the next generation to fully realise their potential. The intergenerational elasticity may also vary across space because of differences in the amount of public- and community-level investments in the human capital of the next generation. Communities with higher than average incomes are more likely to be in a position to make these investments. Similar predictions are implied by more refined models that recognise the distinct developmental stages through which children must pass on their way to developing their full capacities, and the important role families play in this process (Heckman, 2008; Heckman and Mosso, 2014).

While the intergenerational elasticity falls naturally out of the Becker-Tomes framework, it is not the only policy-relevant way of measuring the process. Other perspectives put the focus on group differences and potential nonlinearities. For example, Durlauf and Seshadri (2018) and Durlauf (1996; 2006) develop choice-theoretical models recognising that persistent poverty may be associated not simply with differences in individual investments in human capital but also in social influences on these investments, influences associated with different forms of group membership and identity. Group membership need not have a geographic dimension, but some models associated with the financing of local public goods, like education, certainly suggest that ‘neighbourhoods’ in a geographic sense can be the basis for group identity. This may be the case even if neighbourhood formation is endogenous. Durlauf and Seshadri (2018) explicitly model the influence inequality has on economic opportunity through its impact on socioeconomic segregation. These perspectives give a more refined justification for the possibility that intergenerational mobility may vary across space, and in particular for what Durlauf (2006) refers to as ‘poverty traps’. In this sense, theory is also calling for a certain flexibility in estimation so as to not exclude the possibility of nonlinearities.<sup>1</sup>

If intergenerational cycles of poverty are a concern, then this perspective also calls for appreciating the relevance of both rank-based measures of mobility and indicators highlighting directional movement, in particular the extent of upward mobility from the bottom. Nybom and Stuhler (2017) review alternative rank-based statistics in order to highlight distinct measurement issues: the correlation coefficient, which standardises parent and child incomes for changes in

<sup>1</sup> The theoretical literature has from the start paid, and continues to pay, attention to the possibility of nonlinearities in the intergenerational income elasticity. However, these are associated, in the first instance, with credit constraints among lower-income families, potentially limiting their capacity to finance the human capital development of their children. The implications of this failure in capital markets to permit parents to use the human capital of the next generation as collateral in financing an optimal investment are examined by, among many others, Becker and Tomes (1986), Bratsberg *et al.* (2007), Caucutt and Lochner (2017), Galor and Zeira (1993), Grawe (2004), Hanushek *et al.* (2014), Loury (1981) and Han and Mulligan (2001). Becker *et al.* (2018) develop a model to explain convex non-linearities at the top end of the income distribution without reference to credit constraints.

the degree of inequality over time; the rank correlation coefficient, which depicts the strength of the relationship between income ranks between parents and children; and transition matrices defined according to particular quantiles of the joint distribution of parent and child incomes. Directional movement is also thought of as absolute mobility, reflecting economic growth that determines the extent to which a birth cohort earns more on average than the previous generation. It is also captured by measuring the extent to which any given child rises above, or falls below, the income or the rank his or her parents had in their income distribution (Bhattacharya and Mazumder, 2011; Economic Mobility Project, 2012; Chetty *et al.*, 2017; Ostrovsky, 2017).

The intergenerational elasticity and all these other statistics may be relevant, and the choice between them, as Mazumder (2016) stresses, determined by the analyst's purpose. In a sense, the theoretical literature suggests that it is reasonable to expect mobility to vary with geography, but also that researchers are well-advised to be agnostic in their choice of indicators and careful in interpretation when their purpose is to offer, as is mine, a high-level description.

### 1.2. *The Relevance of the Empirical Literature*

The analysis of the intergenerational transmission of status as measured by earnings and incomes has evolved tremendously, and it is fair to say that recent developments reflect the way progress has generally been made, through a constructive dialogue between economic theory, availability of data and the refinement of appropriate statistical techniques. This is clear in surveys by Björklund and Jäntti (2011), Black and Devereux (2011), Blanden (2013), Corak (2006; 2013), Mulligan (1997) and Solon (1999; 2002). But the American literature is particularly illustrative, with the understanding of intergenerational earnings mobility in the USA being advanced with the release of a series of studies by Raj Chetty, Nathaniel Hendren and their coauthors using American tax-based administrative data. Most notably, Chetty *et al.* (2014) document the extent to which earnings mobility varies across relatively small geographic areas of the USA, these Community Zones being smaller than states but nonetheless larger than metropolitan areas and neighbourhoods. This research is also distinguished by the use of a host of statistics that chart both relative and absolute intergenerational dynamics in ranks.

Chetty *et al.* (2014) put international cross-country comparisons of intergenerational mobility in a different light, suggesting that within-country comparisons are equally important for understanding the process and its policy implications. It may well be that the USA has a lower degree of mobility than many other countries, as suggested by Björklund and Jäntti (2011), Corak (2006; 2013), Jäntti *et al.* (2006) and others, but it is also the case that many regions within the USA are equally, if not more, mobile than many of the most intergenerationally mobile countries, while others are considerably less mobile. It is as relevant, if not more so, for public policy to make comparisons between regions of the same country.

There is a long-standing European literature on intergenerational mobility using administrative data, most of it based in the Nordic countries, with Jäntti *et al.* (2006) illustrating the variety of data sources. Some of this research is subnational and even city-based as in Lindahl *et al.* (2014; 2015), who look at the evolution of intergenerational mobility across four generations in the Swedish city of Malmö.<sup>2</sup> But the encouragement towards regional analysis by Chetty *et al.* (2014) is important for introducing a different flavour to cross-national comparisons.

<sup>2</sup> Barone and Mocetti (2016) also offer a city-based analysis, examining Florence between 1427 and 2011, by using the informational content in surnames.

Jansen (2017) also uses administrative data associated with taxation to study regional differences within the Netherlands, as does Heidrich (2017) for Sweden, and Acciari *et al.* (2019) and Güell *et al.* (2015) for Italy. My analysis is directly informed by this vein in the empirical literature.

### 1.3. *Canadian Studies and the Use of Administrative Data*

Canadian data are well suited for the study of regional differences, but research on intergenerational earnings mobility has focused on estimating the elasticity of father–son incomes at the national level and takes its starting point from Corak and Heisz (1999) and Fortin and Lefebvre (1998). These studies suggest that sons' earnings are higher by about 2% for every 10% increase in fathers' earnings. Fortin and Lefebvre (1998), Corak (2001a) and Chen *et al.* (2017) also offer information on daughters, but it is fair to say that this literature—with the exception of Blanden (2005)—has not stressed mothers and daughters to the degree merited. Corak and Heisz (1999) introduce the use of tax-based administrative data, informally known as the Intergenerational Income Data. These data have been used to estimate various indicators and examine causal mechanisms.<sup>3</sup> But all of this research is conducted at the national level, there being no analysis at finer geography, even at the provincial level. The notable exception is Oreopoulos (2003), who uses the postal codes individuals report on their income tax returns, the T1 forms, to establish their location. He examines the causal impact of neighbourhoods on the long-term outcomes of teenagers growing up in Toronto, focusing on those who lived in social housing projects. I follow this precedent to develop a version of these data appropriate for a regional analysis across the entire country.

This involves updating the data to incorporate both mothers and fathers, and to follow them and their sons and daughters through to 2008.<sup>4</sup> To be more specific, the analysis is based on men and women born between 1963 and 1970. T1 forms first became machine-readable in 1978, and at the time I began this study were available up to 2008, implying that I am able to observe the adult incomes of this cohort of Canadians when they were between 38 and 45 years of age. I link these incomes to the incomes of their parents in the tax years when the children were between the ages of 15 and 19. The sample size is presented in Table 1, showing a weighted total of 3,000,000 parent–child pairs based on an unweighted total of over 2,000,000. The table also shows that the weighted total comes close to the population estimate derived from the Census.<sup>5</sup>

My updating of the data also involves converting the postal codes the parents report on their T1 forms to Census geography codes for the 1986 Census. The analysis focuses upon the Census Division, a geographic unit roughly equivalent to a municipality or county, the legal entity established by legislation and used in many provinces to deliver services. Census Divisions are

<sup>3</sup> See Corak (2001b), Corak *et al.* (2004), Corak and Heisz (1998), Corak *et al.* (2014), Corak and Piraino (2011; 2016), Grawe (2004), Oreopoulos (2003), Oreopoulos *et al.* (2008).

<sup>4</sup> I undertook the direct management of this updating in cooperation with Yuri Ostrovsky, and with the support of Statistics Canada. See the Appendix and Statistics Canada (undated) for more detail on the creation and structure of the data. The original version of these data used by Corak and Heisz (1999) followed sons up to the mid-1990s when they were in their early 30s.

<sup>5</sup> The unweighted sample size represents about two-thirds of the Census population estimate, while the weighted estimates are just over 90% of the Census total. The discrepancy is not simply due to under-coverage of the administrative data associated with the algorithm used to link parent and child income tax returns—which requires that the child have a Social Insurance Number at some point before leaving home—but also to selection rules imposed to ensure that permanent income is accurately measured, an issue I detail in the following sections as well as in the Appendix.

Table 1. *Number of Children Born between 1963 and 1970: Weighted and Unweighted Intergenerational Income Data Sample Sizes Compared to Estimates from the 1986 Census.*

Province/territory	Administrative data		1986 census total	Ratio weighted to census
	Weighted	Unweighted		
Newfoundland and Labrador	84,050	59,000	87,000	0.97
Prince Edward Island	16,750	12,400	18,100	0.93
Nova Scotia	112,900	79,350	121,700	0.93
New Brunswick	91,500	67,600	98,600	0.93
Quebec	796,650	531,000	842,950	0.95
Ontario	1,057,550	796,800	1,191,750	0.89
Manitoba	122,150	91,650	139,500	0.88
Saskatchewan	122,500	81,600	132,700	0.92
Alberta	284,550	184,250	319,550	0.89
British Columbia	304,250	206,200	344,850	0.88
Yukon	2,950	1,700	2,950	1.00
Northwest Territories, Nunavut	7,150	3,600	8,200	0.87
Canada	3,002,950	2,115,150	3,307,900	0.91

Notes: All numbers are rounded to the nearest 50.

Table 2. *Number of Census Divisions by Province and Territory, and Number of Children per Census Division.*

Province/territory	Number of Census		Number of children per Census Division	
	Divisions	Average	Minimum	Maximum
Newfoundland and Labrador	10	8,406	3,946	34,660
Prince Edward Island	3	5,590	2,765	7,714
Nova Scotia	18	6,273	1,173	35,091
New Brunswick	15	6,100	1,539	13,957
Quebec	76	10,482	1,096	180,764
Ontario	49	21,582	1,413	226,374
Manitoba	23	5,310	737	63,250
Saskatchewan	18	6,807	1,841	24,701
Alberta	19	14,976	1,613	95,237
British Columbia	29	10,491	177	121,144
Yukon	1	2,940	2,940	2,940
Northwest Territories, Nunavut	5	1,429	527	3,191
Canada	266	11,289	177	226,374

Notes: Number of children is based on weighted totals per Census Division.

conceptually—but not practically—smaller geographic units than the Commuting Zones used by Chetty *et al.* (2014). Table 2 shows that there are a total of 266 Census Divisions in the country, with 125 of them—almost half—located in the two most populous provinces, Ontario and Quebec. The weighted number of children in each Census Division averages just over 11,000, but varies from a minimum of 177 to a maximum of over 226,000.

## 2. Measuring Intergenerational Mobility

I derive nine statistics of intergenerational mobility in order to measure three different dimensions of equality of opportunity: a concern about income, a concern about rank and a concern about movement through the income distribution, particularly upward mobility from the bottom, but

Table 3. *Summary of Nine Statistics Used to Measure Three Dimensions of Intergenerational Mobility.*

<i>Income mobility</i>		<i>Regression to the mean in parent-child incomes</i>
1. Absolute	$\alpha_j$	least squares estimate of intercept
2. Relative	$\beta_j$	least squares estimate of slope
3. Average income	$\bar{Y}_j$	average permanent income of parents
<i>Rank Mobility</i>		<i>Percentile rank-rank regression</i>
4. Absolute	$a_j$	least squares estimate of intercept
5. Relative	$b_j$	least squares estimate of slope
6. Above parent	$\bar{y}_j$	share with higher income than parents
<i>Directional mobility</i>		<i>Cells of quintile transition matrix</i>
7. Rags to riches	$P_{1,5}$	conditional probability of top income
8. Intergenerational low income	$P_{1,1}$	conditional probability of bottom income
9. Intergenerational privilege	$P_{5,5}$	conditional probability of top income

also intergenerational cycles of poverty and of privilege. Each of these concerns pose particular measurement challenges, and I construct my analytical data to minimise the role of various measurement errors. The nine indicators and the notation used to represent them are summarized in Table 3.

The analysis of income has both an absolute and a relative dimension. The literature focuses attention on the regression to the mean model of incomes across generations:  $\ln Y_{i,t} = \alpha_j + \beta_j \ln Y_{i,t-1} + \varepsilon_{i,j}$ , with  $Y_{i,t}$ , representing the permanent income of a member of generation  $t$  in family  $i$ , and being related log-linearly to  $Y_{i,t-1}$ , the permanent income of a member of generation  $t-1$ . This motivates the use of  $\beta$ , the intergenerational income elasticity, as a relative measure of mobility, indicating the percentage deviation from the mean income in generation  $t$  for a given percentage deviation in generation  $t-1$ , which may vary across communities indexed by  $j$ . The average adult income of children also differs from the average parent income according to the intercept  $\alpha_j$ , a complementary indicator of absolute mobility that may also vary across communities. When the value of the error term is set to its expected value of zero, the anti-log of this equation— $\bar{Y}_t = e^{\alpha_j} \bar{Y}_{t-1}^{\beta_j}$ —shows that the expected income of children from different communities, can vary for at least three statistical reasons: differences in absolute income mobility, differences in relative income mobility and differences in the average incomes of parents across communities.<sup>6</sup> These are the first three of my nine statistics, and for convenience I refer to them respectively as  $\alpha_j$ ,  $\beta_j$  and  $\bar{Y}_j$ . Collectively they focus attention on the income levels of children and their regional variation.

I define ‘income’ as the total income from all sources including both all market sources and government transfers.<sup>7</sup> In addition, my definition refers to the income of both partners, when two

<sup>6</sup> This is true with respect to the geometric mean, but the arithmetic mean of a log-normally distributed variable is  $e^{\mu + \frac{1}{2}\sigma^2}$  where  $\mu$  and  $\sigma$  are the location and scale parameters. Ascribing absolute mobility solely to the intercept term takes liberty with this notion as both the intercept and slope determine the expected value of the child’s adult income. In using this vocabulary the intent is to distinguish the two influences on the child’s outcome, those associated with and those not associated with parental income.

<sup>7</sup> Total income is defined according to the Canada Revenue Agency. From 1982 onwards this refers to: Canada/Quebec Pension Plan Benefits; capital gains/loses; earnings including commissions; interest and investment income; old age security pension; other employment income; other income; pension and superannuation income; rental income; self-employment income (including net business income, net commission income, net farming income, net fishing income, net professional income); employment insurance benefits. Family Allowance benefits are included from 1982 to 1992, and there are also other additions in subsequent years.

are present: any income mothers earn is taken into account to fully represent a family's resources. In both respects, this follows Chetty *et al.* (2014).

The literature has highlighted two sources of potential bias in the estimation of  $\beta$ , an attenuation bias due to measurement error in  $\ln Y_{i,t-1}$  when realised annual income is used to proxy permanent income, and an associated life-cycle bias when income is observed only over a limited part of the life cycle.<sup>8</sup> The attenuation bias arises from mis-measurement of parental permanent income, this is the classical measurement error of a model with only one right-hand side variable. I address this concern by averaging parental income over a five-year horizon, the years the child was 15 to 19 years of age.<sup>9</sup>

The bias due to heterogeneity in life-cycle profiles is particularly pertinent for the accurate measurement of the child's permanent income. Using Swedish data that span practically the entire working lives of parents and children, Nybom and Stuhler (2017) show that estimates of the intergenerational elasticity are in fact very sensitive to a life-cycle bias—both if the children are observed at too young and at too old an age—with the most accurate estimate obtained when the children are between 35 and 40 years of age (Nybom and Stuhler, 2017, Figure 1a). Chen *et al.* (2017) show that life-cycle bias is minimised with the Canadian data I use when the children are in their late 30s to early 40s.<sup>10</sup> I address the possibility of a life-cycle bias by using the average income obtained between 2004 and 2008, when the children are between the ages of 38 and 45 at their oldest, which seems to be a particularly favourable span in the life cycle, annual income measures coming closest to permanent measures. As in the case of parents, my measure of individual income is based on the income of both partners—the child captured in the Intergenerational Income Data and his or her spouse, as identified

<sup>8</sup> Atkinson *et al.* (1983); Grawe (2006); Haider and Solon (2006); Jenkins (1987); Nybom and Stuhler (2017); Solon (1992); Zimmerman (1992).

<sup>9</sup> Measurement error in  $\ln Y_{i,t}$  is not generally seen as generating bias, but rather influencing the efficiency of the estimator. This may be of little concern for a national level analysis using my data because of the large sample sizes, though it may well come into play for subnational estimates, and vary across communities. The use of the term 'permanent income' makes the implicit assumption that parental expectations governing their investment decisions are realised, and accurately captured by some long-run average of actual outcomes. I assume a parent's annual income to be zero if there is no T1 observed in a particular year. The total annual 'family' income is the sum of the father's and mother's incomes if the parents were married or in a common-law relationship in the year the child was linked to them. Only if the mother does not file a T1 in a particular year, or if her Social Insurance Number is unknown, is the family income taken as the father's income. In a similar way the total parental income is the mother's income if the family reports being married or common-law but the father did not file a tax return or if his Social Insurance Number is unknown. In these cases the annual family size is assumed to be one. Single parents, or parents whose marital status changes remain in the sample, and are reflected in changes in the family size. Parental income refers to the combined total income of the mother and the father over the appropriately defined five years divided by 10. This is done regardless of any previous changes in family structure. In cases where the parent was a single parent in the year the child is first linked to the parent, then this family structure is assumed fixed, and that individual's income is divided by five. To be clear, this is a measure of individual permanent income using all income sources, and assuming equal sharing between the adult partners in the household when more than one is present. This average individual income over five years must be at least \$500 to be included in the analytical sample. Preliminary analysis suggested that there may be measurement error at very low incomes, a disproportionate number of individuals often reporting exactly one or exactly two dollars of income on their T1 returns.

<sup>10</sup> This result is for sons, and depicted in their Figure 1. All of this is not to say that life-cycle effects are not also important for parental income. In fact, while I measure child incomes at a given age, parental incomes are measured at a given age of the children, when they are between 15 and 19 years of age. For the most part parents may well be at an appropriate age to capture their permanent income, but this said there is more slippage in this case than with the child adult outcomes. For example, higher-income parents may be older when they have their children. In practice, the distinction does not influence the results, and I prefer grounding parental incomes during the years the children were 15 to 19 because it associates the prevailing parental income to all children at the same stage of their development.

Table 4. *Intergenerational Income Mobility: Absolute Income Mobility, Relative Income Mobility, and Average Parental Community Income.*

Province/territory	Absolute mobility ( $\alpha_j$ )	Relative mobility ( $\beta_j$ )	Average income ( $\bar{Y}_j$ )
Newfoundland and Labrador	8.69	0.180	29,395
Prince Edward Island	8.91	0.159	30,739
Nova Scotia	8.49	0.192	35,158
New Brunswick	8.54	0.189	32,871
Quebec	8.67	0.186	39,700
Ontario	8.67	0.191	44,249
Manitoba	6.98	0.341	36,518
Saskatchewan	8.19	0.238	39,768
Alberta	8.71	0.194	48,544
British Columbia	8.73	0.176	47,185
Yukon	8.62	0.187	42,444
Northwest Territories, Nunavut	8.67	0.175	29,036
Canada	8.52	0.201	42,032

Notes: First two columns are least squares estimates.

from the Social Insurance Number they report on their T1 form for the spouse—when two are present.<sup>11</sup>

Table 4 summarises the provincial and national estimates of these three indicators, the first column offering the absolute income mobility expressed in natural logarithms ( $\alpha_j$ ), the second the intergenerational income elasticity ( $\beta_j$ ) and the third the average income of parents in the community ( $\bar{Y}_j$ ). The standard error for the national estimate of  $\alpha_j$  is 0.005, while for  $\beta_j$  it is 0.0005. The maximum values of these standard errors across the regions are respectively 0.157 and 0.0152, both being for Yukon. The standard errors for the other regions are less than half as large, and for the most part about one-tenth. The implication is that there is considerable variation across the provinces in these parameters.<sup>12</sup>

The second dimension of mobility reflects a concern over position, or ranks. One relevant benchmark from the perspective of parents, and implicitly many policymakers, is whether or not their children will earn more as adults, involving a direct comparison of adult incomes across

<sup>11</sup> A child is deemed ‘married’ if marital status on the tax form is reported as married or common-law. Individuals are considered ‘non-married’ if they are neither married nor living common-law, if the spouse’s Social Insurance Number is unknown, or if they do not have a partner. The total family income is the sum of the child’s income and his or her spouse’s income, or just the child’s income if non-married. As stated, child incomes are calculated between 2004–2008. If in any year the child does not report a T1 form and income is missing, then it is assigned a value of zero. However, the child is required to have at least one T1 during these five years to be included in the sample. The number of dependents in each year is determined by the number of non-zero birth dates for the first six dependents only. The five-year average annual income is computed as the sum of all annual family incomes over the 2004–2008 period divided by the sum of the family sizes over the same period. While my definition of income follows Chetty *et al.* (2014), as does the use of ‘family’ income, there are differences in how the tax systems in the two countries treat common-law relationships. The Canadian tax system is such that it is more appropriate to treat the incomes in a common-law relationship as if the partners are married. All incomes are measured in 2014 constant dollars based upon the Consumer Price Index. As with parents, I keep only parent–child pairs if the child’s average income is not less than \$500.

<sup>12</sup> The intergenerational elasticity for the country as a whole is estimated to be 0.201, lower than the estimates of 0.32 and 0.23 reported respectively for sons and daughters by Chen *et al.* (2017). However, the results are not strictly comparable since the findings in Table 4 are based on sons and daughters without distinction, and also on the incomes of both partners in each generation. I obtain least squares estimates of 0.223 for boys and 0.172 for girls when I define child outcomes in terms of own incomes and undertake separate estimations by gender. This suggests the remaining differences have to do with my use of the total income of both partners as the measure of parental income.

Table 5. *Intergenerational Rank Mobility: Share of Children with Incomes Above their Parents, Absolute Rank Mobility, Relative Rank Mobility, and the Estimated Percentile Rank for a Child Raised by the Average Parents in the Bottom Half of the Income Distribution.*

Province/territory	Above parents ( $\bar{y}_j$ )	At 38 to 45 years of age			At 31 and 32 years of age		
		Absolute ( $a_j$ )	Relative ( $b_j$ )	Expected Rank	Absolute ( $a_j$ )	Relative ( $b_j$ )	Expected Rank
Newfoundland and Labrador	0.724	35.3	0.273	42.1	33.2	0.277	40.1
Prince Edward Island	0.708	35.1	0.245	41.2	35.3	0.239	41.3
Nova Scotia	0.627	32.6	0.251	38.9	32.0	0.249	38.2
New Brunswick	0.658	31.6	0.280	38.6	31.1	0.286	38.2
Quebec	0.636	36.7	0.249	42.9	36.9	0.240	42.9
Ontario	0.624	41.0	0.225	46.6	43.4	0.215	48.8
Manitoba	0.630	31.2	0.325	39.3	29.9	0.320	37.9
Saskatchewan	0.657	41.5	0.226	47.1	37.7	0.236	43.6
Alberta	0.618	44.4	0.206	49.5	41.1	0.203	46.2
British Columbia	0.549	39.6	0.184	44.2	39.9	0.185	44.5
Yukon	0.570	36.3	0.248	42.5	38.5	0.176	42.9
Northwest Territories, Nunavut	0.689	34.1	0.281	41.1	31.4	0.283	38.5
Canada	0.625	38.3	0.242	44.3	38.4	0.240	44.4

Source. Sample shares and least squares estimates and predictions using Statistics Canada, Intergenerational Income Data.

generations within a family. I calculate this indicator as the fraction of children with total incomes higher than their parents. Another reference point is from the child's perspective. In this case, what also matters is the rank in the income distribution of all children, and how tightly it is related to the rank of the parents in their income distribution. The Spearman rank correlation coefficient offers a statistical measure of this type of intergenerational association, say between the percentile rank of children and the percentile rank of their parents. Chetty *et al.* (2014) emphasise this measure, making use of the least squares regression  $y_{i,t} = a_j + b_j y_{i,t-1} + \epsilon_{i,j}$ , where the lowercase letters are used to indicate the child and parent percentile ranks in their respective income distributions, and by implication  $\epsilon_{i,j}$  is uniformly distributed. These two statistics  $a_j$ —absolute rank mobility—and  $b_j$ —relative rank mobility—are, along with the share of children making more than their parents, the next three of my nine measures of mobility.

Measurement error in both parent and child incomes matters when it comes to accurately estimating rank mobility. This is the main reason why I measure child incomes as a five-year average between 2004 and 2008. The general sense in the empirical literature, however, is that life-cycle biases may not be as severe since there may be a tendency for income ranks to be established much earlier in the life cycle, and not change with years in the labour market. This is particularly important in the context of the Chetty *et al.* (2014) study, which is based on adult incomes measured between the ages of 30 and 32, and is likely one of the reasons the authors pay much less attention to income-based statistics. I also define a second analytical sample based on income averaged over the years the child was 31 and 32 years of age in order to assess the robustness of rank-based measures to life-cycle considerations.

Table 5 presents the estimates of rank mobility for the country and each of the provinces and territories. The share of children in each province or territory having more total income than their parents varies from a low of 54.9% to a high of 72.4%. For the country as a whole it is 62.5%. Two sets of estimates are presented for the other indicators. The first is based on the same data as the estimates in Table 4, measuring ranks when the children are in their late 30s and early 40s. The second in the last three columns is based on ranks calculated when the children are 31 and

32 years of age. At the national level there is no significant difference between them. Children raised by parents at the bottom percentile of the income distribution can expect to rank at the 38th percentile of their adult income distribution, a position that they would attain by their early 30s. The difference between a child raised at the very top and one raised at the very bottom is about 24 percentiles in both cases, and the expected rank of a child raised by parents in the bottom half, at the 25th percentile, is the 44th percentile.<sup>13,14</sup>

Figure 1 depicts the national estimates with more nuance by presenting the expected rank of children according to each percentile of parental incomes. The results make clear that the linearity assumption does not fully capture the patterns at the two extremes of the parental distribution. The least squares estimates of the intercept and slope are respectively 38.3 and 0.242, and accurately summarise the data through the bulk of the distribution, from roughly the 10th percentile through to the 80th or possibly the 90th. Children raised by parents ranking below the 10th percentile are expected to attain lower ranks than the linear assumption suggests. The patterns at the top end also show a distinct non-linearity, with rank mobility of children with parents between the 80th and 90th percentiles flat-lining, and then rising for those above the 90th, particularly among those raised by parents in the top three percentiles.

The third dimension of ‘equality of opportunity’ deals with directional mobility, particularly with the capacity of individuals raised by parents in the lower part of the income distribution to move to the top. This focuses on a particular cell of a transition matrix between parent and child ranks:  $P_{o,d}$ , where  $P$  is a transition probability, and  $o, d = 1 \dots Z$  refer to the origin and destination income ranks of the parent–child pair. I set  $Z$  to 5 to focus, for the most part, on the quintile transition matrix. This complements the other indicators because linearity is not assumed and because it speaks to the direction of movement.

Rags to riches mobility may be a particularly pertinent dimension of intergenerational mobility for public policy, but transition matrices involve an adding-up constraint, with both the rows and columns summing to one. As such, the movement from the bottom to the top is related to both intergenerational cycles of poverty—the probability that a child from bottom-income parents will occupy the bottom position in the next generation—and intergenerational cycles of privilege—the probability that being raised by top-income parents predisposes the child to being in the top as an adult. As Corak (2016a), Reeves (2016) and Milanovic (2016) stress, if children cannot escape

<sup>13</sup> I should note that the two sets of columns in the table are not strictly comparable because they differ not just by the age at which child outcomes are measured, but by different underlying samples. The last three columns are based on a subset of the full sample, only those born between 1967 and 1970. Because of a limitation in the construction of the data, I am only able to offer estimates for 31- and 32-year-olds using this younger cohort of individuals, putting aside those who were born between 1963 and 1966. The data has information on the spouses of the sample members beginning only in 1998, and this information is required to derive the ‘family income’. This implies that the sample members born between 1963 and 1966 were 31 between 1994 and 1997, before I can link them to their spouses. Only those born in 1966 are partially in scope, being 32 in 1998, but all those of the younger cohort turn 31 and 32 years of age after the spousal information first becomes available. This said, I also replicate the findings in this table—indeed, all the findings in the paper—using the adult information on the younger cohort. These findings are reported in the Online Appendix to the paper available at [MilesCorak.com/Equality-of-Opportunity](http://MilesCorak.com/Equality-of-Opportunity), and are very similar to the results in the first three columns of Table 5, so this data limitation is of no substantive consequence to the implications to be drawn from the table. It should also be noted that comparing the same cohort at two different ages also implies comparing them at two different times, and time effects may also be confounding the results. This is likely to be more important for absolute income mobility than for rank mobility.

<sup>14</sup> The use of least squares constrains the estimates of the intercept and slope to run the regression line through the midpoint of the distribution leaving only one parameter to be estimated so that  $b = 1 - a/50$ . But this is only the case for the national estimates in the bottom line of Table 5. The estimates for the provinces are not constrained in this way because the data refer to the percentile ranks in the national income distribution, regardless of where the parents and children live. Recall, that children are permitted to be geographically mobile, being ascribed the province indicated by the postal code on their parents’ T1 form when they were linked to them during the teen years.

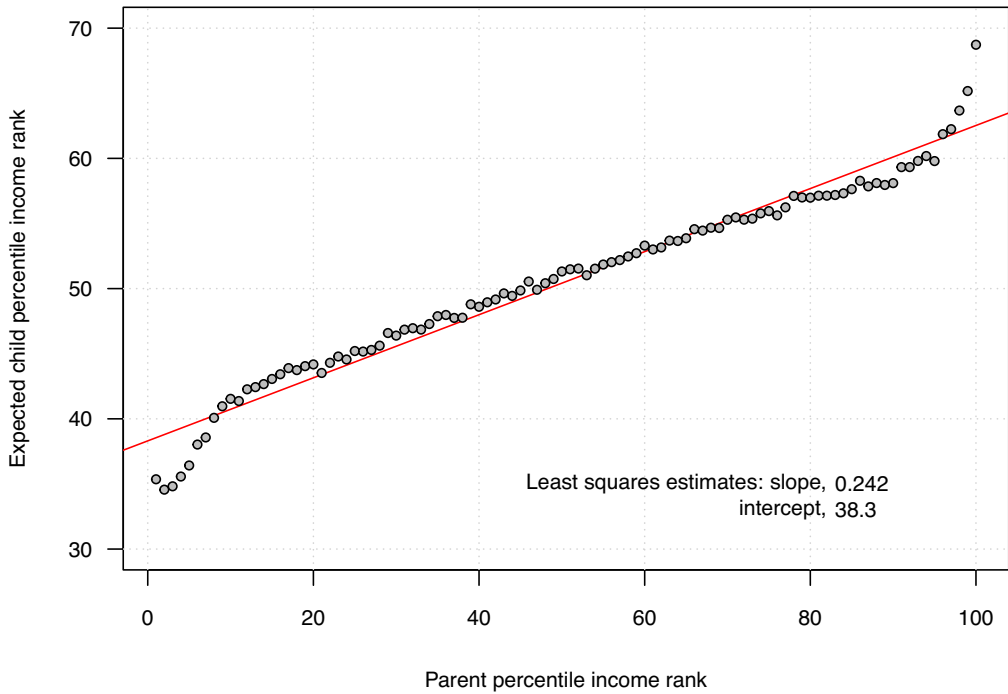


Fig. 1. *Expected Percentile Ranks of Children by Parent Percentile Rank, Canada.*

from the bottom, or if others do not fall out of the top, then the extent of bottom to top movement is limited. In subnational data there is more slack in the adding up constraint between the transition probabilities. All the transition matrices for each Census Division refer to the position of parents and children in the national income distribution. Further, children are geographically mobile and all of the transition probabilities are based on their adult incomes wherever they may be living in the country, it is only the place in which they were raised as teenagers that defines the geography of the analysis. Geographic mobility is an aspect of social mobility.

As such, for any given Census Division, rags to riches movement may not be as tightly constrained by the other cells in the transition matrix. Accordingly, the final three measures of mobility are three cells of the transition matrix, referring to what I will loosely call ‘rags to riches movement’, ‘intergenerational cycles of poverty’ and ‘intergenerational cycles of privilege’. These are operationalised as the three appropriate cells of the quintile transition matrix for each Census Division, though finer matrices and in the extreme a percentile transition matrix can be imagined when sample size permits. They are:

$$P_{1,5} = \Pr\{Y_t \in top | Y_{t-1} \in bottom\}$$

$$P_{1,1} = \Pr\{Y_t \in bottom | Y_{t-1} \in bottom\}$$

$$P_{5,5} = \Pr\{Y_t \in top | Y_{t-1} \in top\}.$$

Nyblom and Stuhler (2017 figure 1[d]) find in their data that transition matrices tend to be robust to life-cycle biases, for example the bottom-to-top quintile probability being accurately estimated

after about the age of 27 or 28. O'Neill *et al.* (2007) also study the statistical properties of transition matrices, stressing the fact that measurement error in both child and parent outcomes need to be addressed to avoid bias. But it should be noted that the measurement error is non-classical. The diagonal elements of the transition matrix will be understated if there is measurement error in the permanent incomes of children. By implication, off-diagonal elements will be overstated, and in particular there will be a tendency to overstate rags to riches movement. The extreme corners of the transition matrix, in particular  $P_{1,1}$  and  $P_{5,5}$ , or  $P_{100,100}$  in the case of a percentile transition matrix, will be at risk of being understated because ranks are bounded from below and above, any measurement error lending bias in only one direction.

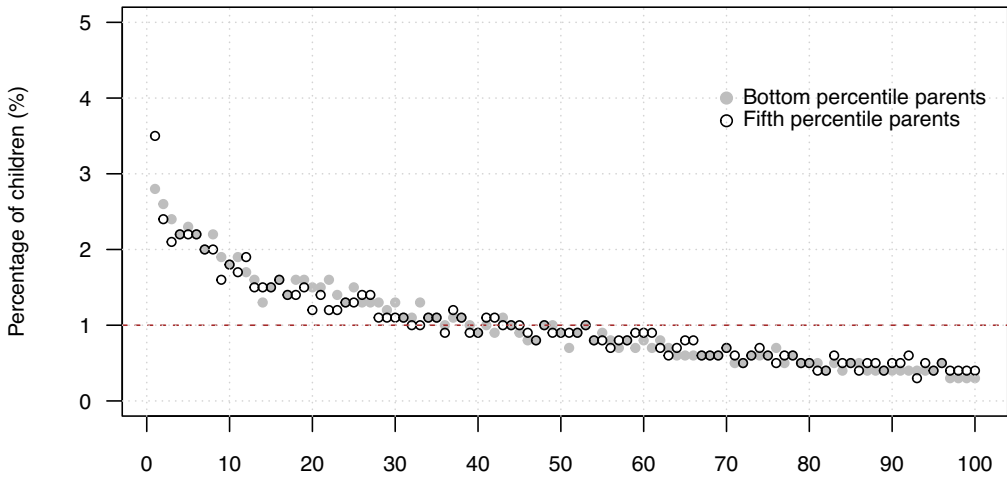
I examine selected parts of the percentile transition matrix in order to assess the risk of bias in the cells of the quintile matrix. The contrast between the percentile rank of children raised by bottom-percentile parents, and those raised by bottom-vingtile parents suggests that this bias is present, but is unlikely to play a role when the data are segmented more widely as quintiles. This is illustrated in the top panel of Figure 2. Children of bottom-percentile parents can fall no lower in their income distribution, but this is not the case for those raised in roughly similar circumstances, whose parents ranked at the fifth percentile. There is a clear gradient in both of these rows of the percentile transition matrix, with each probability of ranking in percentiles one through about 30 being greater than 1%, and the probabilities of rising above the median being lower. The two groups of children share very common probabilities beyond the third percentile, with the major difference between them being in their chance of ranking in the bottom three, and most notably, the very bottom percentile. Children of bottom-vingtile parents have a higher chance of being bottom-percentile adults than their counterparts, and at the same time they also have a slightly lower chance of ranking in the second and third percentiles. The most likely outcome in adulthood of both groups is to be ranked in the bottom percentile, but the chance of this happening for the children of bottom-vingtile parents is notably higher: 3.5% versus 2.8%.

These discrepancies in the lowest ranks may in part reflect non-classical measurement error leading to an underestimation at the very extreme of the percentile transition matrix, namely the probability that children raised by bottom-percentile parents will also be bottom-percentile adults. Mobility out of the bottom for children from the very bottom is possibly understated, while at the same time the adjacent off-diagonal elements, the second and third percentiles, are overstated. The magnitudes suggest that this type of measurement error is likely not to be strongly at play with a wider categorisation, and in particular with quintile transition probabilities.

The transition probabilities at the other extreme of the income distribution are very different. The bottom panel of Figure 2 shows the slice of the transition matrix conditional on having top-income parents. The children of the top 1% are as likely to stay in the top fifth of their income distribution as they are to fall into the bottom 80%. While they face a 49.9% chance of being in the top fifth, their most likely outcome is, at 8.6%, to remain in the top 1%. For similar reasons, we might well expect this to be an understatement.<sup>15</sup>

<sup>15</sup> Nybom and Stuhler (2017) also show that the upward mobility of the children of the very poorest parents tends to be overstated when there is measurement error in both parent and child incomes, almost without regard to whether child incomes are measured at 25, 30 or 40 years of age. The prospects of those raised by parents above the bottom quintile, including those at the very top, tend to be accurately captured if incomes are measured when the children are 40 years of age, though less so at 30 years as there can be considerable overstatement of downward mobility from the very top because top-earners may often have short periods of low income. As the authors point out this is likely for any number of reasons: top earners may spend a period taking more leisure as top incomes are also associated with higher wealth, they could be making money outside the country, or they could be engaging in tax avoidance. This bias is not likely to

(a) Children of bottom one and of bottom fifth percentile parents



(b) Children of top percentile parents

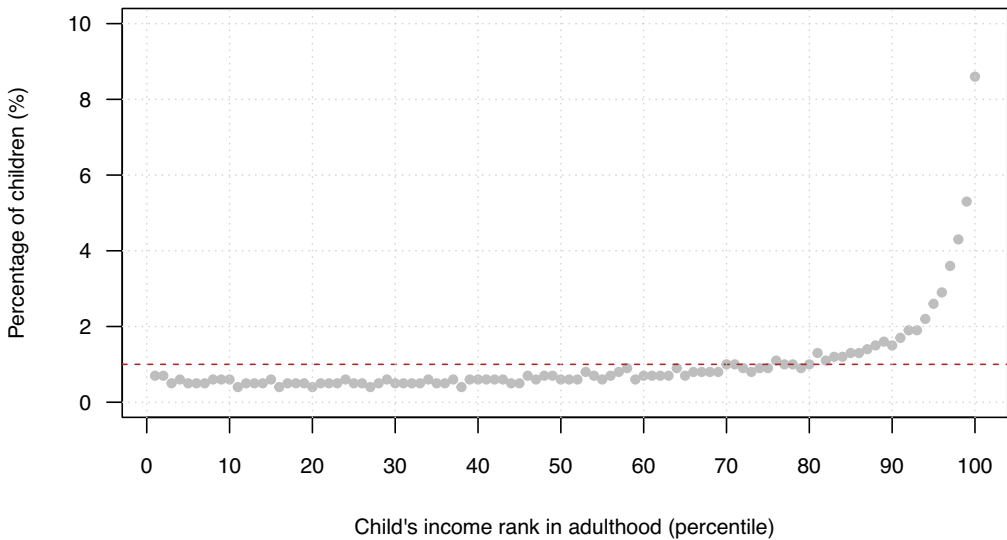


Fig. 2. Potential Measurement Error in Intergenerational Directional Mobility: Percentile Ranks of Children Raised by Bottom-ranking and Top-ranking Parents.

Table 6 offers the three quintile transition probabilities at the heart of my analysis, plus a fourth. The entries in the last column show that the probability a child raised by middle-ranking parents will be a middle-ranked adult hovers around 0.2. If there were no relationship between parent and

apply with equal force for those with low life-time earnings as they are less likely to have an intermittent period of very high income.

Table 6. *Intergenerational Directional Mobility Based on Selected Transition Probabilities: Rags to Riches, Intergenerational Cycle of Poverty, and the Intergenerational Cycle of Privilege.*

Province/territory	Rags to riches ( $P_{1,5}$ )	Cycle of poverty ( $P_{1,1}$ )	Cycle of privilege ( $P_{5,5}$ )	Middle to Middle ( $P_{3,3}$ )
Newfoundland and Labrador	0.087	0.321	0.295	0.210
Prince Edward Island	0.077	0.278	0.279	0.232
Nova Scotia	0.071	0.350	0.256	0.212
New Brunswick	0.061	0.352	0.264	0.214
Quebec	0.091	0.290	0.298	0.233
Ontario	0.141	0.284	0.352	0.210
Manitoba	0.076	0.414	0.296	0.238
Saskatchewan	0.141	0.277	0.333	0.222
Alberta	0.185	0.259	0.375	0.200
British Columbia	0.120	0.298	0.256	0.226
Yukon	0.117	0.371	0.295	0.196
Northwest Territories, Nunavut	0.100	0.397	0.391	0.178
Canada	0.114	0.301	0.323	0.219

Source. Selected entries from the quintile transition matrix calculated using Statistics Canada, Intergenerational Income Data.

child incomes each entry in the quintile transition matrix would be 0.2. A closer look at all the underlying transition probabilities reveals that middle-income parents—pretty well regardless of where they live—are as likely to witness their children fall in the income distribution as they are to witness them rise. This contrasts with the dynamics at the two extremes. Children raised by parents in the bottom quintile stand a 30% chance of also being bottom-quintile adults, and those raised by top-quintile parents have an even slightly higher chance of staying in the top, approaching one-third. There is more variation in these probabilities across the provinces, the intergenerational cycle of low income ranging from as low as 0.26 to over 0.4, and the chances that a child from bottom-quintile parents will rise to the top quintile ranging from 6.1% to almost 19%.

Nybom and Stuhler (2017, page 819) conclude by stating that ‘researchers need to exercise particular caution when studying long-distance mobility, the inheritance of poverty, or the inheritance of top incomes’. These concerns are all addressed by my choice to average both parent and child incomes over multiple years, to measure incomes at an appropriate point in the life cycle, during the late 30s to mid-40s, and also likely helped by the definition of income as encompassing the income of both partners, the use of administrative as opposed to survey data and a focus on transitions across rather wide margins like quintiles.

### 3. The Geography of Intergenerational Mobility

#### 3.1. Correlations among the Indicators

Table 7 presents the Pearson correlation coefficients between the nine mobility statistics. The values indicate both strong and weak correlations, suggesting that the indicators are signalling different aspects of the process. There are some very tight correlations, the value of  $-0.992$  between absolute and relative income mobility being the strongest: Census Divisions with low absolute income mobility tend to have high values for the intergenerational income elasticity.

Table 7. *Pearson Correlation Coefficients between Nine Indicators of Intergenerational Income Mobility across 266 Census Divisions.*

	Income mobility			Rank mobility			Directional mobility		
	Absolute income ( $\alpha_j$ )	Relative income ( $\beta_j$ )	Average income ( $\bar{Y}_j$ )	Absolute rank ( $a_j$ )	Relative rank ( $b_j$ )	Above parents ( $\bar{y}_j$ )	Rags to riches ( $P_{1,5}$ )	Cycle of poverty ( $P_{1,1}$ )	Cycle of privilege ( $P_{5,5}$ )
Absolute income mobility	1.00								
Relative income mobility	-0.992*	1.00							
Average parent income	0.116	-0.057	1.00						
Absolute rank mobility	0.672*	-0.583*	0.457*	1.00					
Relative rank mobility	-0.859*	0.835*	-0.259*	-0.794*	1.00				
Share above parents	0.272*	-0.25*	-0.713*	0.196*	-0.122*	1.00			
Rags to riches	0.283*	-0.198*	0.484*	0.803*	-0.466*	0.016	1.00		
Cycle of poverty	-0.785*	0.712*	-0.288*	-0.876*	0.79*	-0.332*	-0.509*	1.00	
Cycle of privilege	-0.028	0.091	0.301*	0.367*	0.02	0.061	0.603*	-0.116	1.00

Notes: Column entries are Pearson correlation coefficients.

\* indicates the  $p$ -value for a  $t$ -test of the null hypothesis of zero is less than 0.05.

The average income of parents in the communities is more loosely, if at all, correlated with these statistics.

A summary of how these three parameters come together to influence income mobility is offered in Figure 3, which illustrates a strong positive relationship between the average income of parents and the adult income of children across the Census Divisions. The figure also shows that on average children, as adults, earn more than their parents in all but four Census Divisions.<sup>16</sup> The other 262 Census Divisions all lie above the 45° line. Children from these regions earn in adulthood on average \$13,153 more than their parents, reaching a maximum of \$28,384. This said, there is considerable variation in outcomes, with children from some areas earning more than \$20,000 on average than their parents, while in other areas it is less than \$10,000.

It is natural to think of the pattern in Figure 3 as the result of economic growth, with the value of  $\alpha_j$  representing the influence of growth over the span of a generation. This is clearly the case at a national level, but is not simply so in a regional analysis with the possibility of geographic mobility between regions. If  $\alpha_j$  and  $\beta_j$  do not vary across communities, that is if a national level of analysis were appropriate, then regional differences in the adult incomes of children would be the result of differences in the average incomes of parents. But with a regional analysis, these two parameters also play a role. Differences in  $\alpha_j$ , in particular, are not straightforward to interpret. For example, if this parameter is permitted to vary, then the ratio of predicted child adult incomes between any two communities, say  $H$  and  $L$ , is  $(e^{\alpha_H - \alpha_L})(\bar{Y}_H / \bar{Y}_L)^\beta$ , when there is no regional variation in  $\beta$ . In this case  $\alpha_j$  is just as appropriately viewed as an aspect of the human capital of children, that is associated with their capacity to seize economic opportunities across the country. It is important to stress that I follow Chetty *et al.* (2014) in measuring child outcomes in adulthood without regard to where the children live as adults: their geography is determined according to where they lived as young teenagers. While differences in economic growth between regions may be captured by differences in  $\alpha_j$  for those not moving and contribute to the patterns in Figure 3, this cannot be the sole interpretation. Differences in regional growth rates create incentives to move, and communities vary in the extent to which they develop the capacity of children to respond to these incentives, which is certainly an aspect of their human capital. This capacity is not strongly tied to the average parent income in the community, as the

<sup>16</sup> These exceptions are in northern and coastal areas of British Columbia.




Fig. 3. *Average Incomes of Parents and Children in Each of the 266 Census Divisions.*

statistically insignificant correlation between  $\alpha_j$  and  $\bar{Y}_j$  in Table 7 makes clear. Children raised in economically depressed areas may well experience significant absolute income mobility even if there is no growth in the community of their teen years—that is, the community may still be characterised by a high  $\alpha_j$ —if they have a strong tendency to move, or face low fixed costs of moving, to higher income or faster growing regions.

This sheds light on the rather tight correlation between the average community income of parents and the share of children who as adults have a higher income than their parents. The strong negative relationship between these indicators shows that children from communities with higher than average parent income are less likely to exceed their parents. Children from low-income communities are more likely to surpass the income levels of their parents because of the influence of economic growth on their incomes, and because it is more likely that geographic mobility to communities with higher average incomes will lead to higher incomes.

Table 7 also shows that the correlation between relative rank mobility and absolute rank mobility is tight. The indicators vary together in a way that implies the variation in ranks

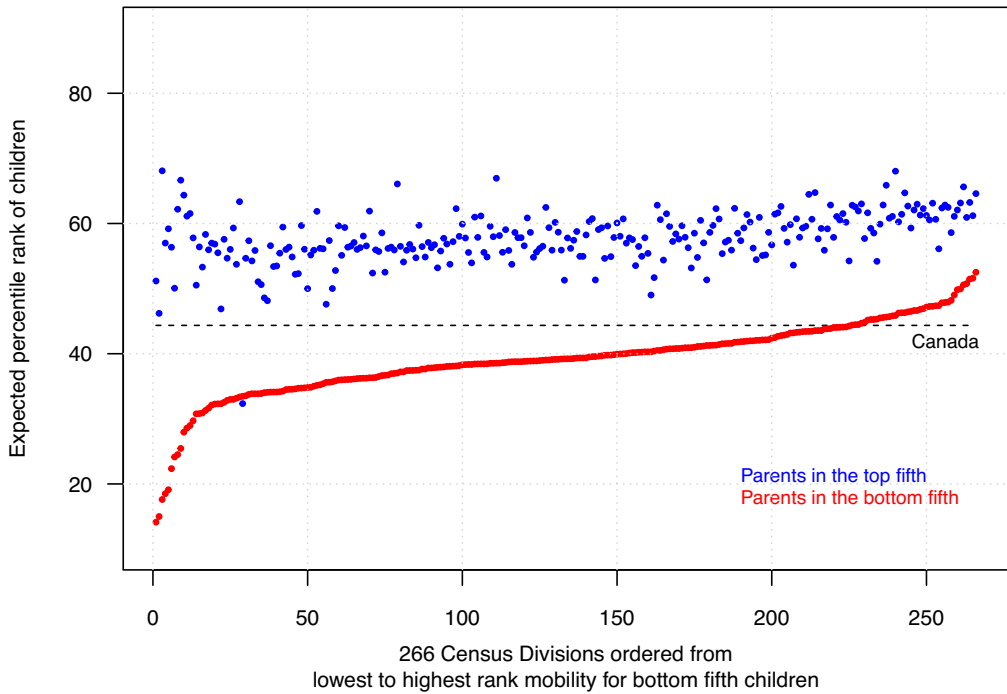


Fig. 4. *Expected Rank Outcomes for Children From Top-fifth and Bottom-fifth Parents by Census Division.*

for children with low-income family backgrounds is greater than for those with high-income backgrounds. For example, the estimated rank mobility coefficients imply that children from top-percentile families will on average rank at the 60th percentile. The standard deviation of these predicted outcomes is 4.3 percentiles. Children from bottom-percentile families can expect to climb to the 37th percentile, the standard deviation across the 266 Census Divisions being higher at 6.3.

The negative correlation between absolute and relative rank mobility is also suggesting that top-ranking parents have a reasonable expectation of seeing their children stay at least in the top half regardless of their location. This is summarised in Figure 4, which is based on the estimates of absolute and relative rank mobility for each Census Division. The least squares estimates of  $a_j$  and  $b_j$  are used to predict the expected percentile ranks of children whose parents were in the top and the bottom fifth of their income distributions, a summary measure of rank mobility that incorporates the influence of both absolute and relative rank mobility. Using the estimates presented in the bottom row of Table 5 suggests that a child raised by the average parent in the bottom half of the income distribution—those ranking at the 25th percentile—would be expected to rise to the 44th percentile, presented as the dashed line in Figure 4. The average rank of children from top-fifth parents is 57.8 with a standard deviation of 4.1, contrasting with 39.1 and 5.8 for bottom-fifth children. The figure shows these predicted chances for each Census Division, ordered according to the expected value for children raised by bottom-fifth parents. In only five Census Divisions is the typical child raised by bottom-quintile parents expected to move into the upper half of the income distribution. There is no strong change in top-quintile expected values according to the ordering of the Census Divisions by expected ranks of bottom-quintile values.

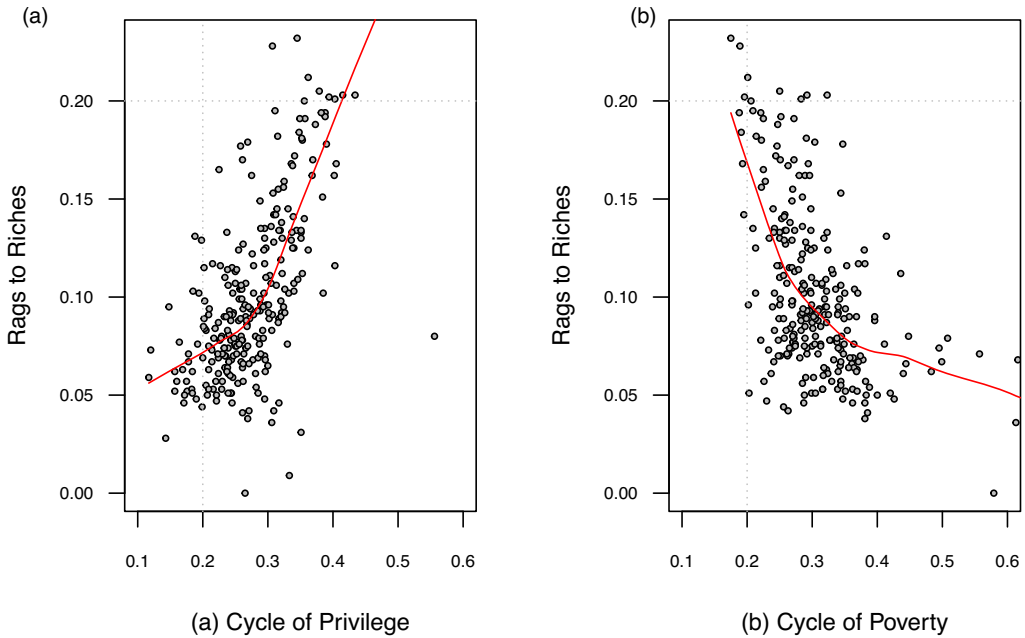


Fig. 5. *Rags to Riches Mobility across Census Divisions and Its Correlation with Intergenerational Cycles of Privilege and Poverty, as Measured with Quintile Transition Probabilities.*

In other words, communities with higher-ranking children from low-ranking families have raised the chances of moving up, and flattened the gradient with parental income in a way that roughly preserves the rank of children from high-ranking families, rather than flattening the slope by lowering the ranks of those highest in the distribution.

This suggests that while high-ranking families are able to promote the relative position of their children regardless of location, this need not limit the upward mobility of children from low-ranking families. The correlations documented in the column of Table 7 labelled  $P_{1,5}$  reinforce this suggestion. A clearer illustration is given in Figure 5, which highlights the positive correlation between rags to riches movement and the cycle of privilege and the negative correlation with the cycle of poverty. The three quintile transition probabilities are graphed in the two panels of the figure for each of the 266 Census Divisions:  $P_{1,5}$  versus  $P_{5,5}$  in panel (a) and versus  $P_{1,1}$  in panel (b).<sup>17</sup> The horizontal and vertical dashed lines are drawn at 0.2 for reference, making clear that it is rare that rags to riches movement is higher than 0.2, while intergenerational cycles of privilege and poverty are almost always above this benchmark. There is a positive relationship between rags to riches movement and cycles of privilege that seems stronger when  $P_{5,5}$  approaches and is higher than 0.3. The negative correlation between  $P_{1,5}$  and  $P_{1,1}$  is much stronger when  $P_{1,1}$  is less than about one-third. As the bottom to bottom transition probability increases from 0.2 to 0.3, the bottom to top probability falls off strongly. Upward mobility from the bottom to the top quintile across the Census Divisions is more a story about the challenges of breaking out of an intergenerational cycle of poverty than it is about breaking into the top. The

<sup>17</sup> Six observations with very high values for  $P_{1,1}$  are omitted from panel (b) of the figure, though they continue to contribute to the estimation of the Lowess line.

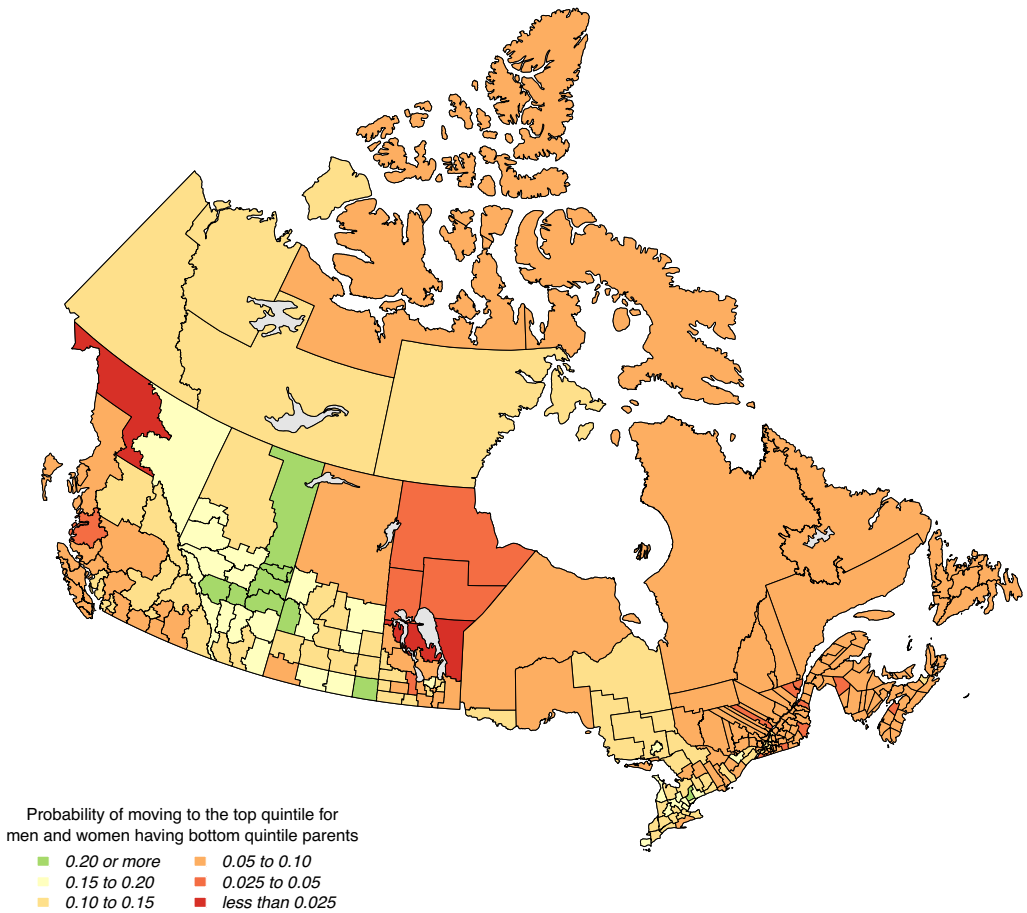


Fig. 6. *Rags to Riches Mobility, Census Divisions Classified According to the Probability that Children of Bottom-Quintile Parents Have Adult Incomes in the Top Quintile.*

very strong correlations in Table 7 between both absolute and relative income mobility and the cycle of poverty, along with their statistically insignificant correlation with the cycle of privilege, reinforce this suggestion.

### 3.2. *Upward Mobility and Cycles of Low Income*

The Canadian landscape of upward mobility is depicted in the map presented as Figure 6, which places each of the 266 Census Divisions into one of six categories according to the probability of moving into the top quintile for children whose parents were in the bottom quintile. Most children—36%—live in the 144 communities where the chances of moving to the top quintile are between 5% and 10%, and a further 34% in the 68 Census Divisions recording a probability of at least 0.10 but not as high as 0.15. There are only two Census Divisions in which the probability of rags to riches movement is less than 2.5%, and a further 14 with chances higher than 2.5%, but not as high as 5%. Together they account for 38,727 children in the weighted sample, or

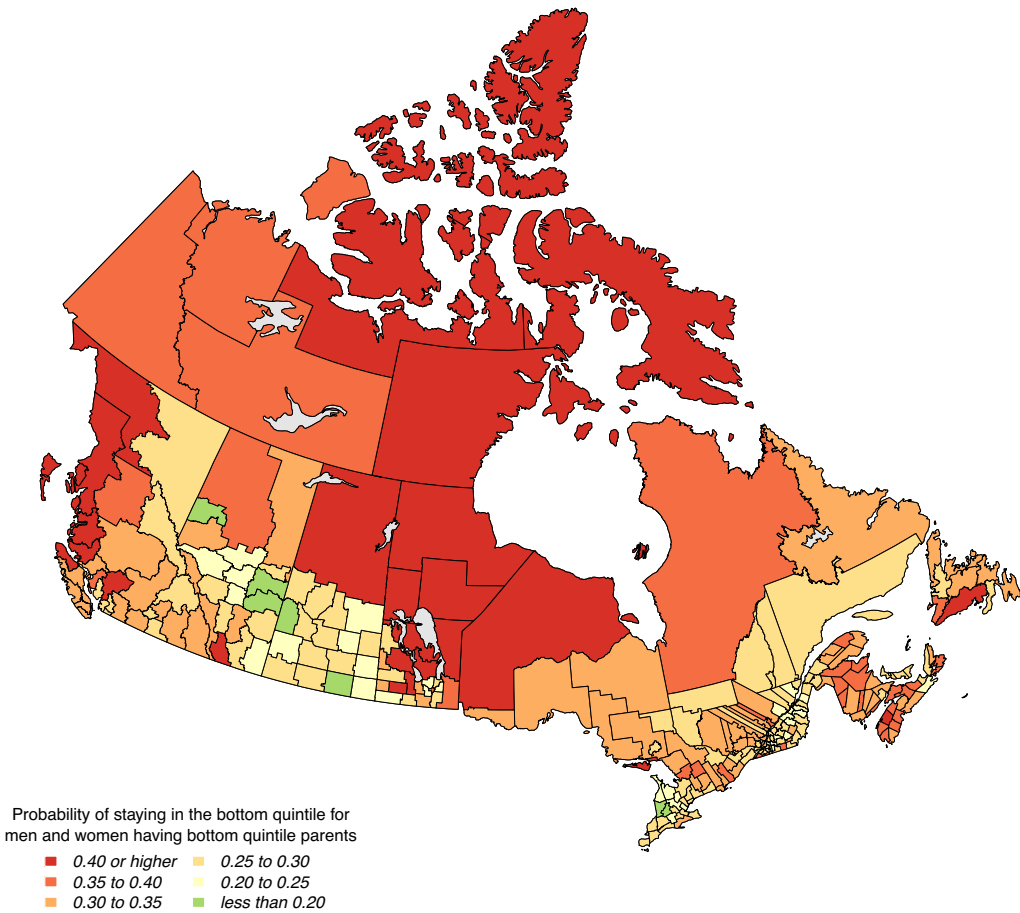


Fig. 7. *Intergenerational Cycles of Low Income, Census Divisions Classified According to the Probability that Children of Bottom-Quintile Parents have Adult Incomes in the Bottom Quintile.*

only about 2.5% of the total. At the other extreme there are nine Census Divisions in which the probability of a rags to riches movement is 0.2 or higher. Only about 3.1% of all children were from these regions. As such, this degree of mobility is a rare feature of the Canadian landscape.<sup>18</sup>

The landscape of intergenerational poverty is mapped in Figure 7, which in a similar manner places Census Divisions into one of six categories according to the bottom to bottom quintile transition probability. These groups range from probabilities of less than 0.20 to more than 0.40. The majority of children—54%—live in the 97 communities where the chances of falling into

<sup>18</sup> Eight of these Census Divisions are in the two high-growth provinces benefiting from a commodity price boom during the time child adult incomes are measured, six in Alberta and an additional two in Saskatchewan. This might be one instance in which the permanent income of children is not measured with complete accuracy, the subsequent commodity price bust probably lowering incomes. The only other Census Division with this exceptionally high rate of bottom to top movement—York Regional Municipality, which is north of and adjacent to what the 1986 Census labels as Toronto Metropolitan Municipality—is the most populous of the group, the others being around a tenth, or even less, in size. This one Census Division accounts for 42.3% of the children who were raised in these high-upward mobility communities.

an intergenerational cycle of low income are between 25% and 30% and a further 24% in the 70 Census Divisions where these chances are at least 0.30 but under 0.35.<sup>19</sup> The strong majority of children raised by lower-income parents face a greater than one-in-four chance of growing up to be relatively lower-income adults, and for many these odds were at least as high as one-in-three. There are 23 Census Divisions with a 40% or greater chance of bottom-quintile to bottom-quintile movement. These communities are all small in population and account for 2% of the total number of children. There are only seven of the 266 Census Divisions in which the probability of a cycle of low income is less than 20%, representing only 1.6% of all children.<sup>20</sup>

### 3.3. *Unsupervised Machine Learning to Cluster Census Divisions*

Together the nine statistics I use offer a broad overview of intergenerational income dynamics, depicting how strongly family and community background determine a child's adult income, where this income ranks in the national income distribution, and the chances of overcoming relative disadvantage and moving from the bottom to the top. As I have suggested, theory offers little guidance on how or even whether we should choose between them, there being little to suggest that one perspective dominates another if our objective is descriptive. And while Tables 4–6 sketch broad geographic differences, getting an overall picture is considerably more complicated at finer disaggregations of the income distribution, particularly for much narrower geographies. The discussion in the previous section implicitly suggests that summarising nine different indicators across hundreds of regions is somewhat more of a challenge. I address this concern by treating the classification of Census Divisions into regions of high or low 'equality of opportunity' as a problem in unsupervised machine learning.

This allows me to cluster the 266 Census Divisions into a limited number of similar regions. Hastie, Tibshirani and Friedman (2009) describe clustering, making clear that the objective of these data segmentation methods is to group units into subsets of greatest similarity according to a metric defined over a set of observed characteristics. I use agglomerative hierarchical algorithms in order to avoid having to make an a priori choice of the number of clusters.<sup>21</sup> In part, I consider the choice of the number of groups to be driven by a communication challenge directed to public policy, hierarchical methods being easily described through a dendrogram that readily offers a visualisation in the form of a tree-like diagram.

My preferred set of results are depicted in Figure 8 and are derived from complete linkage with a Euclidean-based distance over all nine of the intergenerational mobility statistics.<sup>22</sup> The

<sup>19</sup> About 10% of the weighted sample of children are in the 0.20 to 0.25 group, and 7.9% in the 0.35 to 0.40 group.

<sup>20</sup> A common characteristic in these communities is that absolute income mobility is higher than average, and there is a weaker intergenerational elasticity between parent and child incomes. The average parent income in these communities is below the national average. This raises the suggestion that geographic mobility may be an important aspect of intergenerational mobility. The two Ontario communities highlighted with very low cycles of poverty are not areas experiencing significant economic growth, but the distance to nearby regions characterised by high upward mobility—more specifically Toronto—was not great.

<sup>21</sup> To be clear, this is not an exercise in prediction and there is no validation procedure to verify the results. The exercise is descriptive and exploratory. The main class of alternatives would be *K*-means, but I do not use them because they require a pre-determined number of clusters.

<sup>22</sup> Clearly, the results are sensitive to modelling choices and I examined other approaches, with average linkage leading to a broadly similar dendrogram as complete linkage. Single linkage led to different results that could not be judged sensible in the context of my problem. In all cases the intergenerational statistics are scaled. I also used correlation-based distance rather than Euclidean distance. This clusters the 266 Census Divisions into four roughly equally sized groups, but I do not use these results because dissimilarity based on levels is more appropriate for my purpose to describe the country in terms of the degree of mobility.




Fig. 8. Preferred Dendrogram Illustrating the Agglomerative Hierarchical Clustering of 266 Census Divisions, Based on Complete Linkage with Euclidean-based Distance.

vertical distance, not the horizontal distance, between units guides the choice of the number of clusters, the results clearly justifying at least a two-fold division. If the country had to be divided into two groups, into the landscape of ‘us and them’, or another way of putting it, by distinguishing a group that is most sharply different from all the other Census Divisions—in this sense characterised by less equality of opportunity—then the Canadian landscape of economic opportunity would look like the map presented in Figure 9. This is similar to, but not exactly like, the intergenerational cycle of low income depicted in Figure 7, there being only 20 regions of concern.

While this two-fold division of the Canadian landscape is useful as a communication device to point to the first priorities for future research and public policy, it does not represent the most accurate clustering of the Census Divisions. The dendrogram in Figure 8 can be reasonably viewed as suggesting a five-fold clustering: the grouping of Census Divisions highlighted in Figure 9, with the rest of the country divided into four additional groups, using the fourth branch in the dendrogram to discern these clusters.<sup>23</sup> This preferred five-fold grouping is mapped in Figure 10. It offers a more nuanced picture of the landscape of income mobility than any single indicator. This map is not the same as that for the intergenerational cycle of low income, or of rags to riches mobility, depicted in Figures 6 and 7. In Figure 10 many of the eastern regions outside the major urban areas and parts of the north belong to a lower-mobility cluster, labelled as Cluster 4. The more populated areas of Ontario, including all of Southern Ontario are grouped with large parts of Alberta and the southern part of Saskatchewan, and called Cluster 1. This is a high-mobility cluster that also includes the major metropolitan areas of the country.<sup>24</sup>

<sup>23</sup> A sixth cluster consisting of a single Census Division is also indicated by the dendrogram, but it is a region in the very north of the country with a small population, and its selection by the algorithm as a cluster onto itself likely reflects a higher variance among the statistics, rather than any substantive differences.

<sup>24</sup> In particular, Vancouver, some of the municipalities in the vicinity of the island of Montréal, as well as Québec City and its environs, but with the exception of the island of Montréal and the cities in the Atlantic provinces.




Fig. 9. *The Landscape of Us and Them. Census Divisions Clustered into Two Groups Based on Similarities Between All Eight Intergenerational Mobility Measures.*

#### 4. Correlates of Mobility among Census Divisions

Canada ranks as a relatively intergenerationally mobile country when the comparison is to other rich countries, as in the Great Gatsby Curve (Corak, 2013). The positive correlation depicted in the Gatsby Curve—between the intergenerational income elasticity and the Gini coefficient measuring cross-sectional income inequality among parents—should not be given a causal interpretation, but it is predicted by theory and has spurred detailed research on groupings of countries most pertinent for public policy.<sup>25</sup> This said, the Curve privileges one measure of intergenerational mobility—the intergenerational income elasticity—and has not, with the exception of Chetty *et al.* (2014, 1612–14), been examined within countries.

Figure 11 shows a clear positive relationship between income inequality among parents, as measured by the Gini coefficient, and the bottom-quintile to bottom-quintile transition probability,

<sup>25</sup> See Corak (2013; 2016b), Durlauf and Seshadri (2018), Hassler *et al.* (2007), Solon (2004; 2015). Landersø and Heckman (2016) contrast Denmark and the USA, and Bradbury *et al.* (2015) compare the USA to Australia, Canada, and the UK to draw public policy lessons.

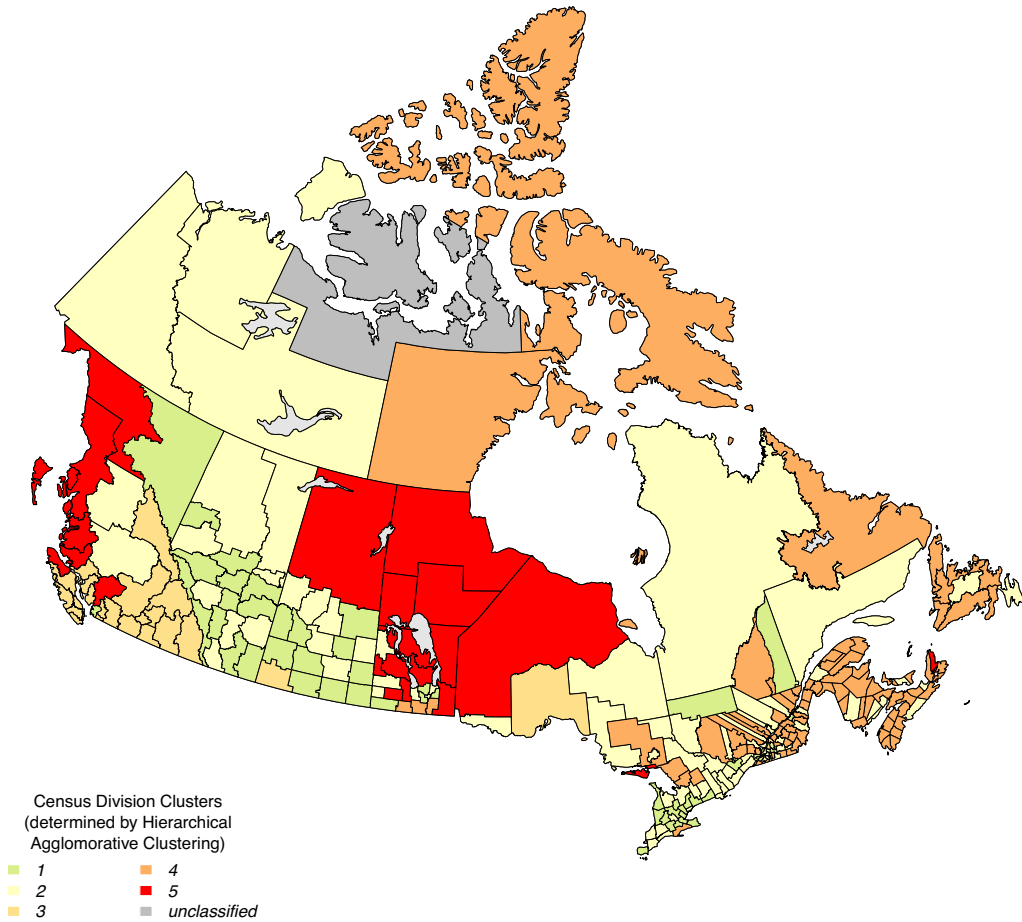


Fig. 10. *The Divided Landscape of Intergenerational Income Mobility, Five Clusters of Census Divisions as Determined by Agglomerative Hierarchical Clustering.*

the unweighted least squares line summarising the positive relationship. Children from low-income families are more likely to become low-income adults when they are raised in regions with higher-income inequality among their parents, a version of the within-Canada Great Gatsby Curve. Bivariate descriptions of this sort are a first step in determining causal relationships. Indeed, the portrait of the degree and nature of intergenerational income mobility across the country in the previous sections immediately raises questions about the factors that characterise high- and low-mobility communities.

Table 8 presents partial correlation coefficients summarising the linear multi-variate relationship between each of the nine statistics measuring mobility, and 17 Census Division characteristics. Some of these characteristics are created with the tax data used in the analysis, others are derived from the 1986 Census.<sup>26</sup> The results are grouped into four broad categories, roughly

<sup>26</sup> The results reported in the table are based upon unweighted multi-variate least squares estimation of data scaled to have mean zero and standard deviation of one. As such they should be interpreted as referring to the standard deviation change in each of the outcomes listed across the columns for a one standard deviation change in the row entries.

Table 8. *Correlates of Intergenerational Income Mobility: Partial Correlation Coefficients Derived from Least Squares Estimation Using Scaled Unweighted Data on Census Divisions.*

Census Division Characteristic	Income mobility			Rank mobility			Directional mobility		
	Absolute income ( $\alpha_j$ )	Relative income ( $\beta_j$ )	Average income ( $\bar{Y}_j$ )	Absolute rank ( $a_j$ )	Relative rank ( $b_j$ )	Above parents ( $\bar{y}_j$ )	Rags to riches ( $P_{1,s}$ )	Cycle of poverty ( $P_{1,1}$ )	Cycle of privilege ( $P_{3,s}$ )
Parents between 25th and 75th percentiles	0.259*	-0.276*	0.009	0.132	-0.36*	-0.096	-0.035	-0.142	0.066
Parent incomes at 50th to 25th percentile	-0.527*	0.545*	0.03	-0.292*	0.509*	-0.137*	-0.063	0.41*	0.143
Parent incomes at 75th to 50th percentile	-0.101	0.061	-0.537*	-0.332	0.13	0.293	-0.321	0.02	0.191
Parent incomes at 90th to 50th percentile	0.243	-0.254	0.309*	0.372	-0.435*	-0.34*	0.342	0.018	0.13
Parent incomes at 99th to 50th percentile	0.024	-0.019	-0.146*	-0.098	0.141	0.294*	-0.02	-0.037	0.214*
Less than high school	0.211*	-0.245*	-0.681*	-0.013	-0.253*	0.537*	-0.161	-0.019	-0.386*
Lone parents	-0.076	0.055	0.062	-0.129	0.069	-0.259*	-0.084	0.189*	-0.143
Not born in Canada	-0.065	0.057	0.231*	0.059	-0.242*	-0.321*	-0.002	0.006	-0.084
English mother tongue	-0.378	0.377	-0.449*	-0.492	0.501	0.247	-0.436	0.117	0.152
French mother tongue	-0.016	0.032	-0.637*	-0.209	0.325	0.735*	-0.152	-0.08	0.472
Parents not married	-0.057	0.011	-0.091	-0.325*	0.166	-0.097	-0.37*	0.231*	-0.121
Population 25 years and older	0.091	-0.067	-0.043	0.161*	-0.014	0.216*	0.229*	-0.113	0.212*
Parents filing in French	-0.258	0.229	0.05	-0.377	0.081	-0.358	-0.678*	0.059	-0.687*
Moved to Census division	0.182*	-0.182*	0.002	0.189*	-0.206*	0.014	0.325*	-0.09	0.176*
Employment in manufacturing	0.005	-0.037	-0.01	-0.158	0.072	-0.116	-0.043	0.105	-0.04
Employment in primary sector	-0.23*	0.249*	0.168*	-0.002	0.204	-0.153	0.159	-0.026	0.096
Employment in top 10 three-digit SICs	0.022	-0.049	-0.009	-0.097	-0.133	-0.114	-0.102	0.156	-0.16

Notes: Column entries are least squares estimates using scaled data on 266 Census Divisions. \* indicates the  $p$ -value for a  $t$ -test of the null hypothesis of zero is less than 0.05.

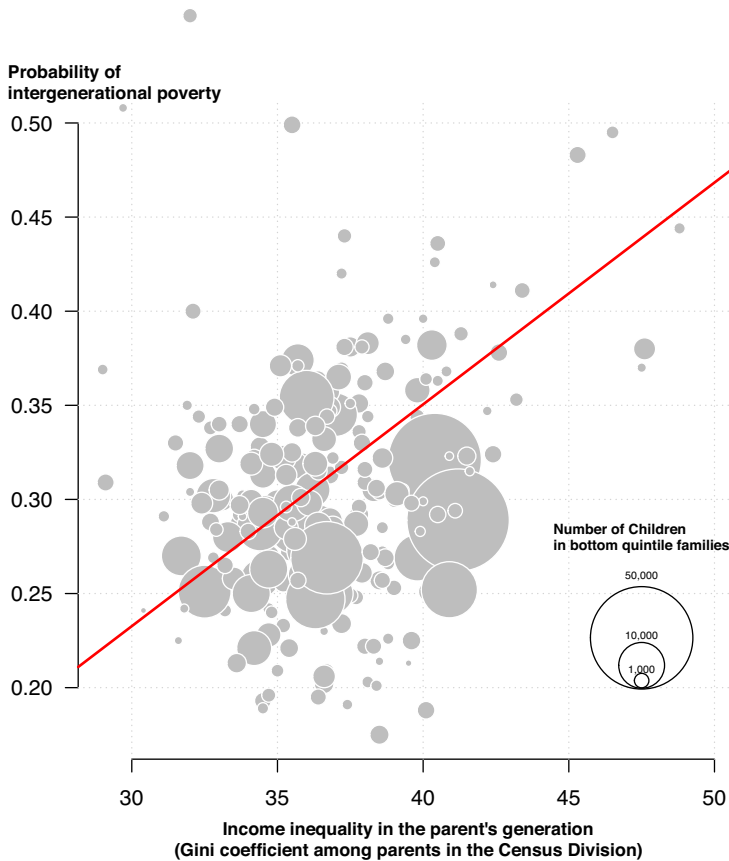


Fig. 11. *The Relationship between the Intergenerational Cycle of Poverty and Income Inequality among Parents.*

corresponding to four important aspects to which theory might be interpreted as drawing attention: income inequality; family and demography; social capital; and the structure of labour markets and opportunities for employment (Corak, 2013).

The first panel highlights the partial correlations with alternative measures of inequality. The Great Gatsby Curve remains evident across the Canadian landscape when inequality is understood to refer to inequality in the lower half of the income distribution. The ratio of incomes between the median income parents and those at the bottom quartile is consistently and strongly associated with the indicators of intergenerational mobility, conditional on all the other variables in the table. The bigger the gap between middle- and bottom-quartile incomes, the lower the amount of absolute income mobility and the higher the intergenerational elasticity. A similar pattern is clear with absolute and relative rank mobility. This type of inequality is also negatively associated with the share of children in the community who will have more income than their parents, and notably it is strongly positively associated with the intergenerational cycle of poverty, though not with the two other indicators of directional mobility. The share of middle-income parents—the fraction of parents between the 25th and 75th percentiles of the national income distribution—plays a

similar, if more attenuated, role in this multi-variate regression, a higher share tending to be associated with more mobility. In contrast, top end income inequality, as measured by the ratio of parent incomes at the 99th percentile to incomes at the median, is not as widely associated with the mobility indicators, though it tends to be positively associated with having more income than parents, and with the intergenerational cycle of privilege.

The next panel groups together a number of demographic and population characteristics. The majority of these are calculated from Census data and refer to population shares in each Census Division: the share of the population 25 and older not having a high school diploma, the share of lone parent households, the share not born in Canada and the share with English and French mother tongue.<sup>27</sup> The most notable partial correlations are between the fraction of single parents in the community and the fraction of children in the analytical sample whose parents were not married on the one hand, and the intergenerational cycle of poverty on the other. Higher population tends to be associated with more absolute rank mobility and more rags to riches mobility as well as cycles of privilege.

Correlations are found with an indicator of geographic mobility, the share of the population in the Census Division who in 1986 reported living in a different Census Division in the 1981 Census. Coleman (1988) suggests that geographic mobility serves as a negative indicator of social capital, the implication is that changing place breaks networks and access to broader community resources. The indicator I offer is imperfect as it is not an individual-level measure of geographic mobility—whether or not the individuals in the analytical file moved—but rather a community level characteristic. A high proportion of residents who moved into the region may reflect the degree to which the region is a pole of growth and economic opportunity, as likely as it is to reflect a negative sense of community and engagement. The correlations seem to suggest the former is more likely, with a higher share of movers associated with greater absolute and relative income, rank mobility and rags to riches movement.

The final block in the table presents indicators of labour market structure, specifically the nature and diversity of employment opportunities during the time parents were raising their children: the proportion of those 25 years and older employed in the manufacturing sector, in the primary sector and the proportion working in the top 10 three-digit industries. The primary share tends to be associated with lower-income mobility, while the manufacturing share and the concentration of employment have no statistically significant correlation with any of the mobility indicators conditional on all the other variables.

All of these findings are descriptive and leave open important questions about the causal mechanisms. How to think about the role of place in the process determining mobility is an open question. Parents living in less rich and more unequal Census Divisions may differ along some unobserved dimensions from parents in other regions even if both groups have the same rank in the national income distribution. In this sense, geography is informative about the parents, a view that contrasts with one ascribing a causal role to place. Oreopoulos (2008) reviews the literature from a Canadian perspective, suggesting that the causal role of place is less important in Canada than in the USA, but this is a topic that should be revisited in light of my findings and data. For example, Chetty and Hendren (2018a,b) use a clear identification strategy based on sibling differences to examine the causal role of place, finding that it is important and offering county

<sup>27</sup> The variables 'Parents not married' and 'Parents filing in French' refer to the information on the tax form of the parents of the children making up the analytical file, while 'Population 25 years and older' is an estimate of the Census Division adult population derived from the Census.

level estimates of the causal impact of neighbourhoods for the USA. Deutscher (2018) uses a similar strategy with Australian data.

## 5. Conclusion

The Canadian geography of economic opportunity is a landscape of considerable mobility, but also considerable diversity. Close to two-thirds of the cohort I examine have attained adult incomes that surpass their parents, but this varies according to the community in which they were raised in a way that reflects both absolute growth in their incomes and how strongly their income relative to the average is tied to their parents' relative income. There is also considerable rank mobility. Upward rank mobility is strong for children whose parents were in the bottom half of the income distribution, though generally not so high as to imply a move into the top half of the income distribution. Overall the typical child raised in the bottom half of the income distribution, by parents who ranked at the 25th percentile, can expect to rise almost 20 percentiles, but this varies depending upon place, amounting to almost nothing in some regions and to as many as 30 percentiles in others. Broadly speaking, upward mobility from the bottom does not seem to be constrained by the possibility that children raised by top-income parents are disproportionately likely to also be top-income adults, but rather by the chances that children raised by bottom-income parents have lower chances of moving out of low income. A significant majority of children of low-income parents live in parts of the country where their chances of growing up to be low-income adults are greater than one-in-four, and for many the chances of this intergenerational cycle of poverty are as high as one-in-three. The Canadian landscape is divided into five non-contiguous areas. Low mobility regions tend to be outside of urban areas, distant from poles of growth, while high mobility regions tend to be urban, or close to major urban areas, but not exclusively so with children growing up in many rural regions being among the most mobile.

In uncovering these descriptive findings I make use of a large bank of administrative data associated with the Canadian income tax system that continues to offer potential to more thoroughly explore these issues. My analysis overcomes a number of measurement problems, but, just as important, it recognises that intergenerational mobility can be defined in a host of different ways. While economic theory offers some guidance for empirical research, the subject of intergenerational mobility is both inherently multi-disciplinary and of ongoing public policy concern, implying that this guidance is less than complete. For this reason I am agnostic with respect to an appropriate metric, and make use of nine related indicators.

My objective is purely descriptive, and for this reason the analytical data set I construct encompasses both men and women, covers the entire country, and uses as broad a measure of income as possible. I uncover some general patterns by documenting the correlations between mobility and a host of Census Division characteristics. Most notably, regions with higher mobility tend to be regions with lower-income inequality, but it is income inequality in the lower half of the income distribution that is most strongly associated with mobility, almost regardless of which indicator is used to measure the child's outcome, but particularly with intergenerational cycles of poverty. In contrast, top end inequality is not strongly associated with intergenerational mobility in the multivariate analysis I offer, an important exception being that it is positively correlated with cycles of privilege among the rich.

At least three avenues of future research suggest themselves. First, attempts at interpreting geographic diversity in intergenerational mobility need to appreciate the nature of economic

growth, its variation across the country and the associated role of geographic mobility. These are aspects of my findings that merit closer attention, and are clearly hinted at by the suggestion that the country contains five broad landscapes of economic opportunity. For example, some parts of southwestern Ontario cluster into a highly mobile landscape in spite of declining fortunes associated with restructuring of the manufacturing sector. This may have something to do with proximity to the city of Toronto, a pole of economic growth. The costs of moving to capture opportunity may be relatively lower than in more outlying areas. My division of the Canadian landscape into a landscape of 'us and them' highlights areas that should be a priority for future research, recognising the need for understanding the roles of both inclusive economic growth and enhanced geographic mobility.

Second, all of the nine measures of mobility I use are based upon 'family' income, referring to the total income of both partners in the household when more than one is present, having the advantage over previous research of including children raised by single parents and more generally mothers and daughters. As a result mobility reflects not just the dynamics of the labour market, but also the marriage market and demographic trends. I find that cycles of poverty and rags to riches movement are correlated with indicators of family structure independent of their association with income inequality. Future research should attempt to parse these influences, examining regional variations across gender in a way that distinguishes labour market incomes and household production. This would highlight the importance of schooling, assortative mating and trends in labour market participation and fertility, and how they vary across socioeconomic groups and across regions.

Third, within country analyses may well add more nuance to cross-country comparisons if mobility is measured in a way that is truly comparable in both income levels and ranks. Chetty *et al.* (2014) offer a framework that hopefully will continue to stimulate similar research in other countries. Their study is clearly the stimulus for my analysis of Canadian data. But as more within-country pictures begin to develop, researchers should not lose sight of the continued importance of cross-national comparisons. The natural comparison of my findings is to the USA, and a reading of Chetty *et al.* (2014) would seem to suggest that Canada is characterised by more rank mobility, and slightly less diversity in outcomes in directional mobility. However, the scope to make informed cross-country comparisons of this sort is only partial because the age cohorts are not the same, because Canadians are not placed in the American income distribution, and because not all of the indicators can be calculated for the two countries. My co-authors and I construct a within and cross-country comparison of these two countries for similar cohorts, and by using percentiles derived from the American income distribution to measure rank mobility, an approach future researchers may wish to consider (Connolly *et al.*, 2019).

At the broadest level my analysis suggests the need for more research to assess the nature of the relationship between income, rank and directional mobility. For example, it is natural for policymakers to give priority to upward mobility among the poor, but this has both an absolute and a relative dimension. Friedman (2006, 95) recognises that these dimensions may well be substitutes in the minds of many citizens, suggesting that 'taking more steps to move a society towards greater fairness or more equal opportunity typically does impose risks, as well as costs, on at least some people. The importance of economic growth for this purpose is that rising incomes make people more willing to accept these risks and costs in the interest of what they take to be a better society for themselves as well as others'. My agnostic approach to the measurement of mobility suggests that empirical research needs more guidance in appreciating the welfare implications of the host of statistical indicators that are used in the literature, and the trade-offs

between them. These implications remain an open issue in the interpretation of my descriptive results.

## Appendix

### *The Creation of the Data and Its Representativeness*

The algorithm used to construct the Intergenerational Income Data (IID) is described in the appendix to Corak and Heisz (1999), and detailed in Statistics Canada (undated). I first developed the IID with colleagues at Statistics Canada in the mid-1990s for the study of intergenerational income dynamics. The data have since been updated in a number of ways, most notably by adding more annual observations on income as they have become available through the tax files submitted to Statistics Canada from the Canada Revenue Agency.

The algorithm creating the file links income tax information (the T1 form) for the members of particular age cohorts through time, and to their parents whose T1 forms are also tracked through time. The data used in this paper are made up of two cohorts of men and women whose T1 forms were linked, mostly through their Social Insurance Numbers, to their parents' T1 forms: those 16 to 19 years old in 1982, and those 16 to 19 years old in 1986. The matching process between parents and children involves up to five years of tax information in order to find a child filing for the first time (and hence having a Social Insurance Number) while still at home according to their tax filing record. For example, the 1982 cohort refers to those individuals who were 16 to 19 years of age in 1982 and who filed a T1 form while at home at some point between 1982 and 1986. In a similar way, the 1986 cohort refers to those in this same age bracket who first filed taxes at some point between 1986 and 1990. It is necessary that individuals file an income tax return while still at home in order to establish a link between their Social Insurance Numbers and the Social Insurance Numbers of their parents. This permits the longitudinal tracking of family members. Individuals who did not file an income tax return while still at home are not included in the IID. Parental T1 forms are available from 1978, though the first possible year that children can be linked to their parents is 1982. Both parents and children are followed through their T1 forms up to 2008, when the children are between the ages of 38 and 45. In addition, the T1 forms of the married and common-law partners of the children are also obtained from the disclosure of their Social Insurance Numbers on the child's adulthood T1 form. The spousal information, however, is available beginning only in 1998. At the time this study began the 2009 income tax data were available only in a preliminary version, and therefore not used.

The coverage of the underlying target populations varies. Detailed examinations based on a comparison with the Census reveals an under-reporting of children with lower-income parents and those living in large metropolitan areas. The former likely reflects the tendency of children from lower-income families to leave home at a younger age, and the latter an under-representation of immigrants who arrived in the country in subsequent years and whose parents came to the country after the mid-1980s. The immigrant under-representation is the natural outcome of using a longitudinal data set, and is not a concern as these cohorts of parents are not likely to have any Canadian tax-based income when the children were growing up. Weights calculated from the Census and additional subnational information on reporting patterns across the income distribution have been derived, and are used throughout the analysis. Past research estimating national level statistics has generally found that it makes little difference whether or not the weights are used. Oreopoulos (2003) pays particular attention to these issues of under-representation in

the context of a study focused on Toronto, and confirms the representativeness of the data. But more caution may be warranted when the objective is to cover the entire country at a relatively fine geography.

### *The Creation of the Weights*

The derivation of the weights is described by Cook and Demnati (2000). The weights are constructed using the 1986 Census and the Longitudinal Administrative Databank (LAD): the former offers an anchor for the population estimates, and the latter allows adjustments by family income. The LAD is a 10% sample of the T1 Family File (T1FF) and includes all children who are dependents of their tax-filing parents regardless of whether they (the children) filed taxes or not. In this sense it is more encompassing than the IID, which is also drawn from the T1FF. However, some children are imputed in the LAD based on other tax and benefit information. This is one reason why it is not an appropriate source for intergenerational mobility studies, along with the fact that it is a one-in-10 sample of T1 filers and offers limited information on the child's gender.

The weights are derived by using only children from the LAD between the ages of 16 and 19 for each cohort, including tax filing and non-filing dependent children. This data allow a comparison of IID under-coverage according to parental income. Eleven income classes based on the parental total market income, in bands of \$10,000 and including \$100,000 and over, are defined for geographic areas determined by the first two digits of the postal code. Some aggregation of adjacent areas was undertaken when the number of observations proved to be very small. This involved the aggregation of 36 areas into 12, and combining the Yukon, Northwest Territories, and Nunavut into one area. This is to say that while the weights are distinguished by parental income, this is done within geographic areas broader than the Census Division.

Weighted counts are derived from the LAD for each cohort, by income class and by these geographic areas. The counts are then used in conjunction with the 1986 Census to adjust for gender, a variable not available in the LAD because some children are imputed. The IID was adjusted with 1986 Census data by age, and for under or over-coverage of the Canadian population by the tax system. This leads to estimates of the number of children by province, age in 1986 and gender, using a weighted 20% file from the 1986 Census. Another set of counts are made from the IID using the basic weight created in the step outlined in the previous paragraph. These counts were made by age in 1986, gender and province. A census adjustment factor is computed using these counts. The final weight is constructed by multiplying the basic weight, from the first step, with the adjustment factor, from the second step. Cook and Demnati (2000) describe the method. The point is that the weighted counts from the IID are meant to be representative of the Canadian population belonging to these cohorts, accounting for gender, geography and parental income.

### *Assigning Census Geographic Information*

The postal code is a six character identifier developed and used by Canada Post Corporation for mail sorting and delivery, and takes the form AnA nAn where A represents an alphabetic character and n is a numeric character between 0 and 9. The first three characters are referred to as the Forward Sortation Area (FSA) and the latter three as the Local Delivery Unit (LDU). The postal code became universal on the T1 files beginning in 1982. The postal code in the year

children are first linked to their parents is used to determine the family's geography, and is the basis for linking Census information to the IID.

Statistics Canada ([undated](#)) describes the process used in converting this information to Census geography codes. The following summary is extracted from this source. All family members should have the same postal code on the assumption that they are co-residents. In fact this is the case for 3,227,271 of the 3,463,712 family records in the full file. Of the 236,441 families reporting at least one different postal code between the members, 190,608 had at least two identical postal codes and this was used as the family postal code. The remaining 45,833 contained 26,876 that had at least two identical Forward Sortation Areas, and this is used as the basis for constructing the family postal code if the LDU was the same across the family members. A similar process was based on the LDU: 6,008 records had at least two family members reporting the same LDU, and this is used as the last three digits if the FSA was reported to be the same across family members. If the postal code was not missing for the remaining 12,949 records, then the postal code of the father was used. If it was missing, that of the mother was used, and finally if that was also missing the postal code reported by the child is used. If all of these codes are missing, the family's postal code was assigned to be missing, and 1,063 records fall into this last category.

Postal codes do not necessarily correspond with the boundaries of the geographic units used by the Census. The Postal Code Conversion File (the GEORES3C Program) is used to append Census geography codes. The complete set of records processed by this program had 3,463,712 records, and of these 30,300 required further processing. In the end only 11 of these could not be assigned a valid geographic code. The Postal Code Conversion File added 1996 Census information to the Family File, and this was converted to 1986 Census information using an auxiliary program (EA96286). Version 3 of this program, written in 1998, was used. For more detail, see the Statistics Canada website: Postal Code Conversion File (PCCF), Reference Guide, 2013. Statistics Canada Catalogue no. 92-154-G, available at <http://www.statcan.gc.ca/pub/92-154-g/92-154-g2013001-eng.pdf> [accessed on May 14, 2015].

### *Census Geographic Coding*

The Census of Canada is organised along three different geographic hierarchies: a national hierarchy covering the entire country; a metropolitan hierarchy that applies to urban areas; and a postal code hierarchy based on the Canada Post Corporation coding. The Enumeration Area is the basic building block that is nested within other broader geographic areas of the national hierarchy. There are 44,042 Enumeration Areas in the 1986 Census. The Census Division (CD) is a geographic unit in the national hierarchy, and refers to legislatively determined areas such as counties, regional districts, regional municipalities or other legislated areas at the provincial level. For the most part their boundaries are established by provincial law for regional planning and the provision of services. These areas are not legislated in four provinces—Newfoundland and Labrador, Manitoba, Saskatchewan, and Alberta—and have been created by Statistics Canada for the purposes of data dissemination. There were 266 CDs in the 1986 Census. The Census geography hierarchies as they existed in 1986 are described in Statistics Canada (1987), which includes a finer geographic area, the Census subdivision. The analysis in this paper uses the national hierarchy at the provincial and Census Division level. Information at the Census subdivision level is also derived. This information is available as a spreadsheet at [MilesCorak.com/Equality-of-Opportunity](http://MilesCorak.com/Equality-of-Opportunity).

*Data Quality Issues*

I create a number of different analytical files that allow me to assess the robustness of the analysis. The first alternative file includes all incomes, rather than using a lower cut-off of \$500 as the basis of selection into the analytical file. The second alternative is based on this cut-off but includes only the 1986 cohort of individuals. I point out below that their geographic locations as teenagers might be considered to be more accurately determined. In similar ways, alternative files are also developed for boys and girls separately. The core of the analysis, that reported in the text, is based on the file that uses all cohorts. This maximises the sample size, and maintains an income cut-off of \$500 because this is a better approximation of permanent income, avoiding measurement error in the lower end of the distribution and permitting the derivation of logarithms used in the assessment of income mobility. I also create similar analytical files from the 1986 cohort that is based on the children being 31 and 32 years of age. These alternatives allow me to assess measurement issues associated with life-cycle biases, and make comparisons to a similarly aged cohort used by Chetty *et al.* (2014).

As alluded to, the IID weights will not perfectly correspond to the subnational geography at which the analysis is actually conducted. The weights are based on province and broad subprovincial regions defined by the first two digits of the postal code. As a result, there is no guarantee that weighted population counts from the IID will correspond with Census counts for areas as small as the Census subdivision. For this reason the analysis has been conducted at a level no finer than the Census Division, and with both weighted and unweighted data. This paper reports only the results from the weighted data, the choice not making substantive differences.

Figure A1 illustrates an important quality concern by presenting a very loosely defined estimate of the ‘coverage ratio’ for each Census Division, based on the analytical file used in the analysis. This is the ratio of the population estimate in the tax-based analytical files—unweighted and weighted—to the population estimate for this age cohort of Canadians as derived from the 1986 Census. These should be loosely interpreted. For example, it is not clear that these ratios should be equal to one since the analytical file is based on a number of selection rules to ensure that the individuals are still alive and present in Canada in the five year period ending in 2008, and also that a reasonable estimate is made of their permanent incomes. The unweighted coverage ratios range from as low as 0.3 to as high as 0.85, and on average are 0.64. The weighted ratios are, as the Figure also shows, higher, and on average are equal to one. But many are greater than one, though most range between 0.75 and 1.25.

There are two reasons these ratios deviate from one. To some degree it reflects the slippage between the way in which the weights are calculated according to the first two digits of the postal code, while the placement of children into particular Census Divisions is done according to the full six-digit postal code. There are nine Census Divisions in which the ratio of population weighted counts from the tax data to the Census data are less than 0.75, and 11 with ratios greater than 1.25. The results for these communities should be interpreted with caution, though they selectively continue to be included in some of the results reported in this paper. In particular, they continue, with one exception, to be included in the cluster analysis, where there is no explicit reporting or use of any one mobility statistic. All of these communities are very small in population for the age group of individuals that is the focus of attention.

The second reason for the coverage ratios being very low or very high is related to how accurately children may be placed in particular Census Divisions. This has to do with the child’s age in the year linked to the parents. Recall that children are assigned to a geographic area




Fig. A1. Coverage Ratios of the Population of Children in Each of the 266 Census Divisions: Ratio of Sample Population in Tax Data Relative to Census Estimate, Weighted and Unweighted.

according to the postal code on the T1 files in the year they are linked by the IID algorithm to their parents. As suggested, at one extreme this involves using individuals who are 16 to 19 years of age in 1986, and linked to their parents in that year. Situating these children geographically with 1986 Census information is probably an accurate link to the communities in which they were growing up between the ages of 15 to 19, when parental income is measured. But at the other extreme the linkage algorithm involves a cohort who were 16 to 19 years of age in 1982, and some being as old as 23 when linked with their parents in 1986 and when the Census geography is also established. Some fraction of these children and their parents may well have been living in a different Census Division when the children were between 15 and 19, which would date from 1982 and earlier. This raises a question as to how accurately the region in which they actually lived is captured by ascribing them to a Census Division in 1986. Figure A2 illustrates by replicating Figure A1 for the youngest cohort in the IID, those who were 16 to 19 in 1986. The weighted coverage ratios are much more likely to be less than one. There are 31 of the 266 Census Divisions with a coverage ratio greater than one, and only two with a ratio greater than 1.25.

However, the differences between the two samples described in Figures A1 and A2 do not lead to substantive differences in the findings. When the results for all cohorts presented in the body of the paper are compared to those from the subset belonging to the 1986 cohort, there are differences, but generally only to the extent that might be expected from chance. Figure A3 offers one illustration. The data in this Figure contrast least squares estimates of relative rank

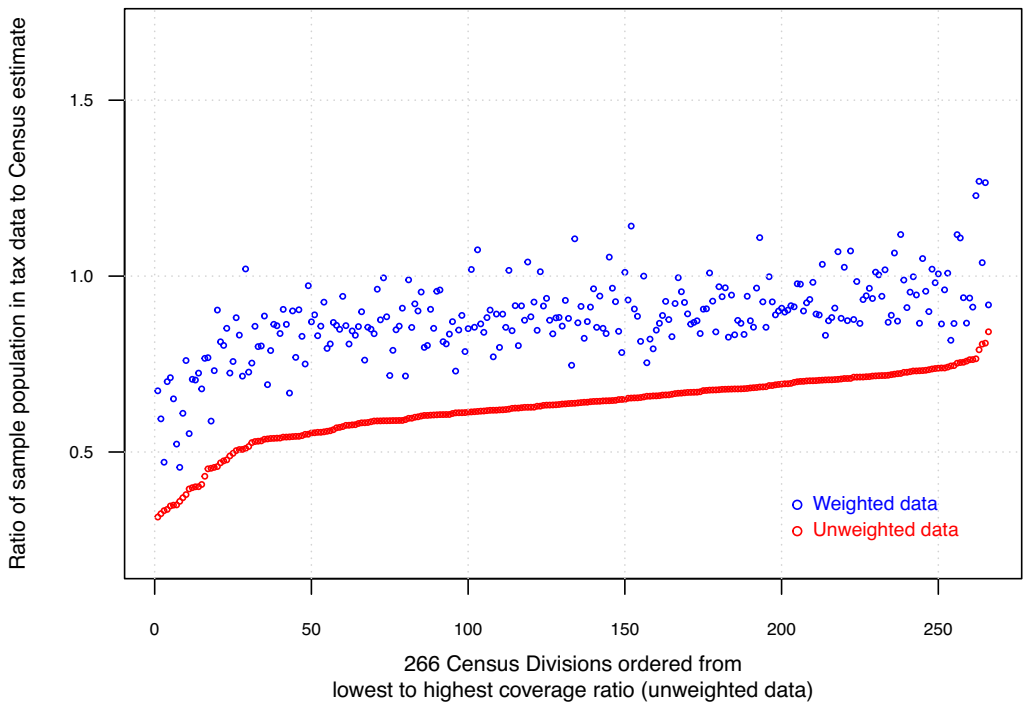


Fig. A2. Coverage Ratios of the Population of Children 16 to 19 Years of Age in 1986 in Each of 266 Census Divisions: Ratio of Sample Population in Tax Data Relative to Census Estimate, Weighted and Unweighted.

mobility— $100 \times b_j$ —between the two samples used in Figures A1 and A2. The vertical lines represent the 95% confidence interval around the estimated  $b_j$  for the full sample. The estimates for the 1986 cohort and their confidence intervals are presented only for those Census Divisions in which the null hypothesis of no difference between the estimates can be rejected. Relative rank mobility is estimated to be statistically different—as determined by a  $z$ -test at the 95% confidence level—between these two samples in 14 of the 266 Census Divisions. This is not greater than what would be expected by chance. This said, when the estimates are found to be statistically different, they tend to be higher using the 1986 cohort in all but two cases.

An Online Appendix available at <https://milesorak.com/equality-of-opportunity/> replicates all the tables and figures in the original working paper version of this paper using the younger cohort. The differences do not overturn the major conclusions. This said, there are some differences in the classification of Census Divisions into the particular groupings used to create the maps and some tables. These groupings are not based on any notion of statistical significance, so slight differences in the values of the statistics around the cut-offs defining the groupings may lead to different classifications. The Online Appendix also separately compares each of the mobility statistics used in the clustering. In general, there is no systematic difference between the two sets of statistics with the exception of the estimates for absolute income mobility and the average parent income in the Census Division. Both these statistics show a tendency to be higher when the full sample is used. This may also influence the classifications resulting from the clustering

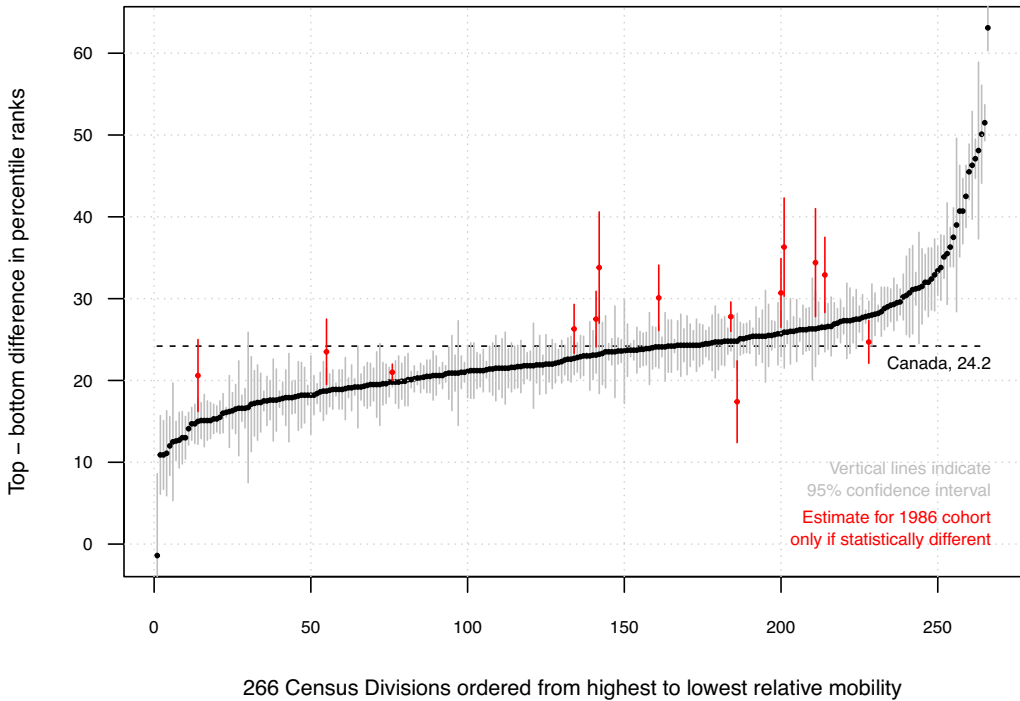


Fig. A3. *Least Squares Estimates of Relative Rank Mobility for all Children and for the Subset of Those 16 to 19 Years of Age in 1986 in Each of the 266 Census Divisions.*

algorithm. However, in some sense this result is to be expected given that the children and their parents have a tendency to be older in the full sample and, as a result, are likely to be earning slightly higher incomes.

*The Graduate Center, City University of New York, NY 10016, USA*

Additional Supporting Information may be found in the online version of this article:

### Replication Package

### References

- Acciari, P., Polo, A. and Violante, G.L. (2019). 'And yet, it moves: intergenerational mobility in Italy', NBER Working Paper Series No. 25732, Mass: Cambridge.
- Atkinson, A.B., Maynard, A. and Trinder, C. (1983). *Parents and Children: Incomes in Two Generations*, London: Heinemann Educational Books.
- Barone, G. and Mocetti, S. (2016). 'Intergenerational mobility in the very long run: Florence 1427–2011', Working Paper 1060, Banca d'Italia.
- Becker, G.S., Kominers, S.D., Murphy, K.M. and Spenkuch, J.L. (2018). 'A theory of intergenerational mobility', *Journal of Political Economy*, vol. 126(S1), pp. S7–25.
- Becker, G.S. and Tomes, N. (1979). 'An equilibrium theory of the distribution of income and intergenerational mobility', *Journal of Political Economy*, vol. 87(6), pp. 1153–89.
- Becker, G.S. and Tomes, N. (1986). 'Human capital and the rise and fall of families', *Journal of Labor Economics*, vol. 4(3, Part 2), pp. S1–39.

- Bhattacharya, D. and Mazumder, B. (2011). 'A nonparametric analysis of black–white differences in intergenerational income mobility in the United States', *Quantitative Economics*, vol. 2(3), pp. 335–79.
- Björklund, A. and Jäntti, M. (2011). 'Intergenerational income mobility and the role of family background', in (W. Salverda, B. Nolan and T. Smeeding, eds.), *The Oxford Handbook of Economic Inequality*, Oxford: Oxford University Press.
- Black, S.E. and Devereux, P.J. (2011). 'Recent developments in intergenerational mobility', in (D. Card and O. Ashenfelter, eds.), *Handbook of Labor Economics*, vol. 4, pp. 1487–541, Part B, chap. 16, Amsterdam: Elsevier.
- Blanden, J. (2005). 'Love and money: intergenerational mobility and marital matching on parental income', *Research Paper No. 272*. Ottawa: Statistics Canada.
- Blanden, J. (2013). 'Cross-country differences in intergenerational earnings mobility: a comparison of approaches from economics and sociology', *Journal of Economic Surveys*, vol. 27(1), pp. 38–73.
- Bradbury, B., Corak, M., Waldfogel, J. and Washbrook, E. (2015). *Too Many Children Left Behind: The U.S. Achievement Gap in Comparative Perspective*, New York: Russell Sage Foundation.
- Bratsberg, B., Røed, K., Raaum, O., Naylor, R., Jäntti, M., Eriksson, T. and Österbacka, E. (2007). 'Nonlinearities in intergenerational earnings mobility: Consequences for cross-country comparisons', *ECONOMIC JOURNAL*, vol. 117(519), pp. C72–92.
- Carney, M. (2014). 'Inclusive capitalism: creating a sense of the systemic', Speech given by Mark Carney, Governor of the Bank of England at the Conference on Inclusive Capitalism, London.
- Caucutt, E. and Lochner, L. (2017). 'Early and late human capital investments, borrowing constraints, and the family', Working Paper 2017-040, Chicago: Human Capital & Economic Opportunity Global Working Group.
- Chen, W.H., Ostrovsky, Y. and Piraino, P. (2017). 'Lifecycle variation, errors-in-variables bias and nonlinearities in intergenerational income transmission: new evidence from Canada', *Labour Economics*, vol. 44(1), pp. 1–12.
- Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R. and Narang, J. (2017). 'The fading American dream: trends in absolute income mobility since 1940', *Science*, vol. 356(6336), pp. 398–406.
- Chetty, R. and Hendren, N. (2018a). 'The impacts of neighborhoods on intergenerational mobility I: childhood exposure effects', *Quarterly Journal of Economics*, vol. 133(3), pp. 1107–62.
- Chetty, R. and Hendren, N. (2018b). 'The impacts of neighborhoods on intergenerational mobility II: county-level estimates', *Quarterly Journal of Economics*, vol. 133(3), pp. 1163–228.
- Chetty, R., Hendren, N., Kline, P. and Saez, E. (2014). 'Where is the land of opportunity? The geography of intergenerational mobility in the United States', *Quarterly Journal of Economics*, vol. 129(4), pp. 1553–623.
- Coleman, J.S. (1988). 'Social capital in the creation of human capital', *American Journal of Sociology*, vol. 94(Supplement), pp. S95–120.
- Connolly, M., Corak, M. and Haecck, C. (2019). 'Intergenerational income mobility between and within Canada and the United States', *Journal of Labor Economics*, vol. 37(S2), pp. S595–641.
- Cook, K. and Demnati, A. (2000). *Weighting the intergenerational income data file*, Ottawa: Statistics Canada, Social Surveys Methods Division.
- Corak, M. (2001a). 'Are the kids all right? Intergenerational mobility and child well-being in Canada', in (K. Banting, A. Sharpe and F. St-Hilaire, eds.), *Review of Economic Performance and Social Progress*, Montreal and Ottawa: Institute for Research on Public Policy and Centre for the Study of Living Standards.
- Corak, M. (2001b). 'Death and divorce: the long-term consequences of parental loss on adolescents', *Journal of Labor Economics*, vol. 19(3), pp. 682–715.
- Corak, M. (2006). 'Do poor children become poor adults? Lessons from a cross country comparison of generational earnings mobility', in (I.J. Creedy and K. Kalb, eds.), *Research on Economic Inequality*, vol. 13, Amsterdam: Elsevier.
- Corak, M. (2013). 'Income inequality, equality of opportunity, and intergenerational mobility', *Journal of Economic Perspectives*, vol. 27(3), pp. 79–102.
- Corak, M. (2016a). 'How much social mobility? More, but not without other things', in (M.R. Strain, ed.), *The US Labor Market: Questions and Challenges for Public Policy*, Washington, DC: American Enterprise Institute.
- Corak, M. (2016b). 'How the Great Gatsby Curve got its name', <https://milesorak.com/2016/12/04/how-the-great-gatsby-by-curve-got-its-name/>.
- Corak, M., Gustafsson, B. and Osterberg, T. (2004). 'Intergenerational influences on the receipt of unemployment insurance in Canada and Sweden', in (M. Corak, ed.), *Generational Income Mobility in North America and Europe*, Cambridge: Cambridge University Press.
- Corak, M. and Heisz, A. (1998). 'How to get ahead in life: some correlates of intergenerational income mobility in Canada', in (M. Corak, ed.), *Labour Markets, Social Institutions, and the Future of Canada's Children*, Catalogue No. 89-553, Ottawa: Statistics Canada.
- Corak, M. and Heisz, A. (1999). 'The intergenerational earnings and income mobility of Canadian men: evidence from longitudinal income tax data', *Journal of Human Resources*, vol. 34(3), pp. 504–33.
- Corak, M., Lindquist, M.J. and Mazumder, B. (2014). 'A comparison of upward and downward intergenerational mobility in Canada, Sweden and the United States', *Labour Economics*, vol. 30, pp. 185–200.
- Corak, M. and Piraino, P. (2011). 'The intergenerational transmission of employers', *Journal of Labor Economics*, vol. 29(1), pp. 37–68.

- Corak, M. and Piraino, P. (2016). 'The inheritance of employers and nonlinearities in intergenerational earnings mobility', in (K. Basu and J. Stiglitz, eds.), *Inequality and Growth: Patterns and Policy*, London: Palgrave Macmillan.
- Deutscher, N. (2018). 'Place, jobs, peers and the importance of teenage years: exposure effects and intergenerational mobility', *Crawford School of Public Policy*, Australian National University, Unpublished manuscript.
- Durlauf, S.N. (1996). 'A theory of persistent income inequality', *Journal of Economic Growth*, vol. 1(1), pp. 75–93.
- Durlauf, S.N. (2006). 'Groups, social influences, and inequality', in (S. Bowles, S.N. Durlauf and K. Hoff, eds.), *Poverty Traps*, pp. 141–75, New York and Princeton N.J.: Russell Sage Foundation and Princeton University Press.
- Durlauf, S.N. and Seshadri, A. (2018). 'Understanding the Great Gatsby Curve', in (M.S. Eichenbaum and J. Parker, eds.), *NBER Macroeconomics Annual 2017*, vol. 32, chap. 4, Chicago: University of Chicago Press.
- Economic Mobility Project. (2012). *Pursuing the American dream: economic mobility across generations*, Washington, DC: Pew Charitable Trusts.
- Fortin, N. and Lefebvre, S. (1998). 'Intergenerational income mobility in Canada', in (M. Corak, ed.), *Labour Markets, Social Institutions, and the Future of Canada's Children*, Catalogue No. 89-553, Ottawa: Statistics Canada.
- Friedman, B.M. (2006). *The Moral Consequences of Economic Growth*, New York: Alfred A. Knopf.
- Furman, J. (2016). 'Equality and efficiency: a global perspective', *Remarks by the Chairman, Council of Economic Advisors, to the World Bank Group Macroeconomics and Fiscal Management Global Practice Annual Forum*, Washington DC.
- Galor, O. and Zeira, J. (1993). 'Income distribution and macroeconomics', *Review of Economic Studies*, vol. 60(1), pp. 35–52.
- Grawe, N.D. (2004). 'Reconsidering the use of nonlinearities in intergenerational earnings mobility as a test for credit constraints', *Journal of Human Resources*, vol. 39(3), pp. 813–27.
- Grawe, N.D. (2006). 'Lifecycle bias in estimates of intergenerational earnings persistence', *Labour Economics*, vol. 13(5), pp. 551–70.
- Güell, M., Pellizzari, M., Pica, G. and Mora, J.V.R. (2018). 'Correlating social mobility and economic outcomes', *ECONOMIC JOURNAL*, vol. 128 (612), pp. F353–403.
- Haider, S. and Solon, G. (2006). 'Life-cycle variation in the association between current and lifetime earnings', *American Economic Review*, vol. 96(4), pp. 1308–20.
- Han, S. and Mulligan, C.B. (2001). 'Human capital, heterogeneity and estimated degrees of intergenerational mobility', *ECONOMIC JOURNAL*, vol. 111(470), pp. 207–43.
- Hanushek, E.A., Leung, C.K.Y. and Yilmaz, K. (2014). 'Borrowing constraints, college aid, and intergenerational mobility', *Journal of Human Capital*, vol. 8(1), pp. 1–41.
- Hassler, J., Rodríguez Mora, J.V. and Zeira, J. (2007). 'Inequality and mobility', *Journal of Economic Growth*, vol. 12(3), pp. 235–59.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., New York: Springer.
- Heckman, J.J. (2008). 'Schools, skills, and synapses', *Economic Inquiry*, vol. 46(3), pp. 289–324.
- Heckman, J.J. and Mosso, S. (2014). 'The economics of human development and social mobility', NBER Working Paper Series No. 19925, Mass: Cambridge.
- Heidrich, S. (2017). 'Intergenerational mobility in Sweden: a regional perspective', *Journal of Population Economics*, vol. 30(4), pp. 1241–80.
- Jansen, M. (2017). *Intergenerational Economic Mobility in the Netherlands*, Master's thesis, Department of Economics, Erasmus University Rotterdam.
- Jäntti, M., Bratsberg, B., Røed, K., Raaum, O., Naylor, R., Österbacka, E., Björklund, A. and Eriksson, T. (2006). 'American exceptionalism in a new light: a comparison of intergenerational earnings mobility in the nordic countries, the United Kingdom and the United States', *Institute for the Study of Labor*.
- Jenkins, S. (1987). 'Snapshots versus movies', *European Economic Review*, vol. 31(5), pp. 1149–58.
- Krueger, A.B. (2012). 'The rise and consequences of inequality in the United States', *Remarks by the Chairman, Council of Economic Advisors, to the Center for American Progress in Washington, DC*, January 12th. [https://obamawhitehouse.archives.gov/sites/default/files/krueger\\_cap\\_speech\\_final\\_remarks.pdf](https://obamawhitehouse.archives.gov/sites/default/files/krueger_cap_speech_final_remarks.pdf).
- Landersø, R. and Heckman, J.J. (2016). 'The Scandinavian fantasy: the sources of intergenerational mobility in Denmark and the U.S.', *Scandinavian Journal of Economics*, vol. 118(4), pp. 794–846.
- Lindahl, M., Palme, M., Sandgren-Massih, S. and Sjögren, A. (2014). 'A test of the Becker-Tomes model of human capital transmission using microdata on four generations', *Journal of Human Capital*, vol. 8(1), pp. 80–96.
- Lindahl, M., Palme, M., Sandgren-Massih, S. and Sjögren, A. (2015). 'Long-term intergenerational persistence of human capital: an empirical analysis of four generations', *Journal of Human Resources*, vol. 50(1), pp. 1–33.
- Loury, G.C. (1981). 'Intergenerational transfers and the distribution of earnings', *Econometrica*, vol. 49(4), pp. 843–67.
- Mazumder, B. (2016). 'Estimating the intergenerational elasticity and rank association in the United States: Overcoming the current limitations of tax data', in (L. Cappellari, S.W. Polachek and K. Tatsiramos, eds.), *Inequality: Causes and Consequences*, vol. 43, pp. 83–129, *Research in Labor Economics*, Emerald Group Publishing.
- Milanovic, B. (2016). 'The downside of upward mobility', <https://glineq.blogspot.ca/2016/09/the-downside-of-upward-mobility.html>.
- Mulligan, C.B. (1997). *Parental Priorities and Economic Inequality*, Chicago: University of Chicago Press.

- Nybohm, M. and Stuhler, J. (2017). 'Biases in standard measures of intergenerational income dependence', *Journal of Human Resources*, vol. 52(3), pp. 800–25.
- O'Neill, D., Sweetman, O. and Van de gaer, D. (2007). 'The effects of measurement error and omitted variables when using transition matrices to measure intergenerational mobility', *Journal of Economic Inequality*, vol. 5(2), pp. 159–78.
- Oreopoulos, P. (2003). 'The long-run consequences of living in a poor neighborhood', *Quarterly Journal of Economics*, vol. 118(4), pp. 1533–75.
- Oreopoulos, P. (2008). 'Neighbourhood effects in Canada: a critique', *Canadian Public Policy*, vol. 34(2), pp. 237–58.
- Oreopoulos, P., Page, M. and Stevens, A.H. (2008). 'The intergenerational effects of worker displacement', *Journal of Labor Economics*, vol. 26(3), pp. 455–83.
- Ostrovsky, Y. (2017). 'Doing as well as one's parents? Tracking recent changes in absolute income mobility in Canada', *Economic Insights*, Catalogue No. 11-626-X, Ottawa: Statistics Canada.
- Reeves, R.V. (2016). *Dream Hoarders: How the American Upper Middle Class is Leaving Everyone Else in the Dust, Why That is a Problem, and What to Do About it*, Washington DC: Brookings Institution Press.
- Solon, G. (1992). 'Intergenerational income mobility in the United States', *American Economic Review*, vol. 82(3), pp. 393–408.
- Solon, G. (1999). 'Intergenerational mobility in the labor market', in (O.C. Ashenfelter and D. Card, eds.), *Handbook of Labor Economics*, vol. 3A, Amsterdam: Elsevier Science.
- Solon, G. (2002). 'Cross-country differences in intergenerational earnings mobility', *Journal of Economic Perspectives*, vol. 16(3), pp. 59–66.
- Solon, G. (2004). 'A model of intergenerational mobility variation over time and place', in (M. Corak, ed.), *Generational Income Mobility in North America and Europe*, Cambridge: Cambridge University Press.
- Solon, G. (2018). 'What do we know so far about multigenerational mobility?', *ECONOMIC JOURNAL*, vol. 128(612), pp. F340–52.
- Statistics Canada. (1987). '1986 census of Canada—geography tape file', Catalogue number CS92-F0023/1987-PDF, Ottawa: Statistics Canada. <http://publications.gc.ca/site/eng/438508/publication.html>.
- Statistics Canada. (undated). 'Intergenerational income data', Family & Labour Studies Division; unpublished mimeo.
- White House, Office of the Press Secretary. (2013). 'Remarks by the president on economic mobility', Speech by President Obama at THEARC, Washington DC, December 4. <https://www.whitehouse.gov/the-press-office/2013/12/04/remarks-president-economic-mobility>.
- Yellen, J.L. (2014). 'Perspectives on inequality and opportunity from the survey of consumer finances', Speech by Chair Janet L. Yellen at the Conference on Economic Opportunity and Inequality, Boston, Massachusetts: Federal Reserve Bank of Boston. October 17th. <http://www.federalreserve.gov/newsevents/speech/yellen20141017a.htm>.
- Zimmerman, D.J. (1992). 'Regression toward mediocrity in economic stature', *American Economic Review*, vol. 82(3), pp. 409–29.