Parental Income in the Labor Market *

Jean-William Laliberté Alexander Whalley

December 2024

Abstract

Are children of high-income families more likely to work at better-paying firms, and if so, why? To answer these questions, we use Canadian administrative data to construct an employee-employer-parent-child matched dataset, which we link to detailed educational records. We use these data to quantify the role of observable human capital (education) and social connections for firm sorting. We find that, in an accounting sense, access to high-paying employers explains roughly half of the transmission of income across generations, as measured by the income rank-rank relationship. Probing mechanisms contributing to the sorting of high-income children to high-paying employers, we find the explanatory power of education largely exceeds that of social connections. We further provide suggestive evidence that some children of high-income families receive preferential treatment when they work at firms their parents own, but the quantitative importance of this phenomenon for overall intergenerational income mobility remains limited.

Keywords:
JEL Classifications:

*Laliberté: University of Calgary and IZA, jeanwilliam.lalibert@ucalgary.ca; Whalley: University of Calgary, alexander.whalley@ucalgary.ca. We are grateful to Nathanael Hammond for excellent research assistance. We also thank Raj Chetty, Henrik Kleven, David Margolis, Olivier Marie, Sacha Kapoor, Matt Staiger, and numerous seminar and conference participants for helpful conversations.

1 Introduction

Income inequality persists across generations. Traditionally, studies of intergenerational income mobility have mainly focused on differences in human capital that can be traced back to differences in either inherited advantages or parental investments in childhood.[1] Yet, firm-specific hiring and compensation policies also contribute significantly to earnings inequality (Song et al., 2019; Card et al., 2013; Abowd et al., 1999).[2] In this paper we ask, are children of high-income families more likely to work at better-paying firms, and if so, why?

One possibility is that children from privileged backgrounds acquire different levels and types of human capital, which gives them access to jobs at better-paying firms. Indeed, prior work has found positive selection into high-pay firms on the basis of skills (Håkanson et al., 2021; Card et al., 2012; Hellerstein and Neumark, 2008). Alternatively, children from different backgrounds may have different social networks (Chetty et al., 2022). Since a considerable fraction of jobs are found through family and friends (Dustmann et al., 2016), social connections may exacerbate inequality in labor market opportunities (Eliason et al., 2023; Staiger, 2022; Bolte et al., 2020; San, 2022; Kramarz and Skans, 2014).

Researchers interested in these questions face important data challenges. For instance, matched employee-employer datasets do not typically contain information on family relationships. Also, administrative tax data rarely include information on educational attainment or other forms of human capital. We overcome these challenges using newly linked data from Canada. Our analytical dataset is based on administrative tax data that include both employee-employer and child-parent linkages, which we match with post-secondary education enrolment and graduation data. Crucially, these highly granular education records include information on the exact institution and program of study.[3] This feature serves two purposes. First, it allows us to examine the role of education with great flexibility, notably permitting for idiosyncratic match effects between firms and specific education programs. For examples, some programs may tailor their curriculum to the needs of local employers, producing specific education-to-job pipelines. Second, we can identify the impact of parental connections on hiring probabilities by comparing people who graduated from the same program and therefore faced similar labor market prospects. Such narrow comparisons help isolate

[1]Examples include work on the role of parenting practices (Doepke et al., 2019), neighborhoods (Chyn and Katz, 2021), colleges (Chetty et al., 2020), and pre-birth factors such as genes (Björklund et al., 2006).

[2]Firms have been shown to contribute to group differences by ethnicity (Gerard et al., 2021), gender (Card et al., 2016), and immigration status (Dostie et al., 2023).

[3]For example, we know if a child completed a Masters in Computer Science from the University of Waterloo or a certificate in Pastry Arts at Vancouver Community College.

the influence of connections on the likelihood that parents and children work for the same employer from other mechanisms such as parents and children making similar educational investments or living in similar areas.

The starting point of our investigation is the intergenerational rank-rank income relationship (Chetty et al., 2014), which effectively partitions children into 100 mutually exclusive parental income groups. To examine the role of firms for mobility, we gradually peal off different layers of mechanisms related to employers and contributing to the persistence of income across generations.[4] The first step is a simple Kitagawa-Oaxaca-Blinder-type decomposition of the rank-rank relationship into within- ("unexplained") and between-firm ("explained") components. We find that, in an accounting sense, the between-firm sorting component accounts for roughly half of the intergenerational income rank-rank relationship. This component conflates differences in firm pay premiums (i.e. firm valued-added) and selection effects due to unobserved heterogeneity across firms. To separate these two channels, we make use of AKM (Abowd et al., 1999) firm effects to account for unobserved heterogeneity in time-invariant earnings potential. We find an important role for firm premiums: they explain close to a third of the transmission of income across generations.[5]

We then seek to understand the mechanisms behind the between-firm sorting component. Formally, we produce counterfactual distributions of children across firms using information on education and parental connections to employers. That is, we predict the probability that a child matches with a specific employer given their education program and whether they are socially connected to that employer (e.g. if one of their parents work there). To do so, we use the dyadic setup of Kramarz and Skans (2014), where the unit of analysis is a worker-firm dyad and the sample includes all possible worker-firm pairs.

The richness of the educational data allows us to account for broad patterns of hiring on the basis of educational attainment as well as match effects between specific institutions and firms. The latter may be important if children of high-income parents are more likely to pursue programs of studies that systematically stream its graduates into high-paying firms. For example, prior work has shown that the proximity between universities and high-wage firms (Weinstein, 2022), as well as employers' on-campus recruiting strategies (Weinstein, 2018) do affect hiring outcomes. Regarding social connections, it is well-established that

[4] Throughout the paper, we use the words "firm" and "employer" interchangeably, although the term employer might be more appropriate since we include both private- and public-sector organizations.

[5] There is a parallel with traditional Kitagawa-Oaxaca-Blinder-type decompositions, where the explained component (due to differences in explanatory variables) can be evaluated at either group's compensation structure (the coefficients on these variables). This formulation is very general, and allows one to plug-in any compensation structure. We follow this logic and plug in AKM effects to isolate the role of firm value-added.

many adult children work at firms that also hire one of their parents (Staiger, 2022; Corak and Piraino, 2011). Using data on business ownership, we also show that parent business owners hiring their own children is a very common phenomenon among families in the top 1 percentile of the income distribution. We therefore consider these two types of social connections separately in our analyses. Importantly, the dyadic setup allows us to quantify the importance of unequal labor market *opportunities* – that is, not just of "realized" matches that involve a parental connections (i.e. cases where a worker does work at their parent's firm), but also of connections to firms that may be part of a worker's outside option.

We find that parental connections are a significant determinant of which firm children match with. The likelihood of working at a given firm increases by 2 to 6 percentage points when one's parent works there. The effect size is a full order of magnitude larger in cases when a parent is the owner of the firm. The quantitative importance of connections for income mobility depends not just on effect sizes but also on the distribution of connections across families. Overall, we find that parental connections explain 10% of differences in the firm sorting component across parental income groups, with the lion share being driven by connections to a parent's employer. Still, the explanatory power of education remains more important, explaining 40% of the between-firm sorting component.

Since parental connections may not only affect which employer someone matches with, but also the remuneration they obtain from that employer, in our last analysis we examine within-firm compensation gaps. Here, we find that children working for an employer that also employs their parents do not earn much more than their unconnected co-workers. In contrast, children of firm owners who work for their parents do earn significantly more than their colleagues. This may be because they receive preferential treatment, or because they are a particularly good fit for the firm. For instance, children of business owners may know the industry very well, having observed their parents navigate this environment growing up. Using firm balance sheet data, we find that firm revenue per worker decreases following the hiring of the owner's child. This piece of evidence lends support to the idea that business owners treat their children differently than they treat other employees. Nevertheless, the overall explanatory power of parental connections for within-firm income gaps is quite small (4%). In comparison, differences in education explain 30% of the extent to which children from high-income families earn more than their co-workers.

Overall, we document an important role of employers for intergenerational income mobility. Parental connections contribute significantly to the persistence of income across generations, and most of it is driven by access to higher-paying firms rather than through within-firm differences in remuneration between connected and unconnected co-workers. Still, the explanatory power of education remains larger in magnitude. Differences in education contribute substantially to both sorting to employers and differences in remuneration among co-workers. This implies a non-negligible fraction of returns to education operate via firm sorting.

We see our work as contributing to several strands of literature. First, we contribute to a large literature on intergenerational inequality (Becker and Tomes, 1979; Solon, 1999; Björklund et al., 2006; Sacerdote, 2007; Chetty et al., 2014; Stuhler, 2018) by documenting an important role for firms. In particular, our work is related to a smaller literature on the role of labor markets and early career shocks as contributors to the persistence of income across generations (e.g. Kaila et al. (2022)). The papers most closely related to ours are Engzell and Wilmers (2022), which uses Swedish data, and Dobbin and Zohar (2023), which uses Israeli data. Both of these papers decompose the intergenerational income elasticity (IGE) into individual and firm components using a AKM-type two-way fixed effects specification. In contrast, we opt for a non-parametric rank-rank approach, which has several methodological advantages. First, it allows the explanatory power of firms (and of the mechanisms we study) to vary flexibly across the parental income distribution. This is particularly important given salient non-linearities at the tails. Second, by effectively partitioning the sample into mutually exclusive groups based on parental income, we can compute counterfactual distributions of workers across firms allowing for education-firm match effects and worker-firm specific relationships based on parental connections. This approach allows for much richer patterns of sorting into firms than alternative methods do.[6] Third, the rank-rank approach allows for the inclusion of individuals with zero income, and is also less sensitive to age-at-measurement and the length of the window one uses to measure income (Deutscher and Mazumder, 2020; Chetty et al., 2014). Our work also innovates in relation to other contemporary work by measuring education at greater levels of detail and by incorporating information on parents' firm ownership.

Second, we contribute to the literature on the importance of social connections in the labor market. Numerous papers have examined the role of employee referrals (Burks et al., 2015; Dustmann et al., 2016), while others have used either parental networks or co-worker based networks to evaluate the value of connections for job finding (Glitz, 2017) and wage bargaining (Caldwell and Harmon, 2019). While important contributions identify the causal

[6] In comparison, evaluations of mechanisms based on the IGE are often done by including individual characteristics as linear controls in a regression of children's average firm premium on their father's log income.

effects of social connections (Eliason et al., 2023; Kramarz and Skans, 2014), we know less about how quantitatively significant social connections are for intergenerational inequality. We contribute by showing that while social connections do contribute to the sorting of workers across firms, their quantitative significance for income mobility is limited.

Finally, we contribute to the literature on education and sorting in the labor market. In a broad sense, different firms and industries hire workers with different levels of education (Card et al., 2024; Engbom and Moser, 2017). But educational institutions can play an important role in the allocation of workers to firms. For instance, Oyer and Schaefer (2016) demonstrate that lawyers segregate into firms on the basis of the law school they graduated from, and Zimmerman (2019) find that graduates of elite colleges in Chile disproportionately occupy leadership positions at publicly traded firms. Rivera (2011) provide qualitative evidence that employers pay substantial attention to applicants' schools, and dramatically favor those from super-elite schools. Relative to these papers showing the importance of university-firm connections, we quantify its total impact on firm sorting by parental income.

2 Data

We use administrative tax data from the Canadian Employer Employee Dynamics Database (CEEDD), linked with the Postsecondary Student Information System (PSIS). The CEEDD covers years 2001 to 2018, inclusive, whereas the PSIS covers years 2009 onward. On the worker side, the CEEDD includes individual (T1) and family (T1FF) tax forms for the universe of tax filers in Canada. On the firm side, it includes balance sheet data (T2 corporate tax returns), and, for private incorporated businesses, a list of all owners and their corresponding ownership shares (T2S50 forms).[7] The database includes all employers in Canada, both in the public and private sectors. Within the public sector, separate organizations (e.g., different ministries, school boards or public utilities) are assigned unique identifiers.

The CEEDD also includes job-level information based on T4 slips and Records of Employment (ROE). This allows us to link workers to firms by finding the employers who issued their T4 slips. Since workers can receive more than one T4 in a year (e.g. if they switch employers in the middle of the tax year), we match workers with the employer from which they earned the most within the year.

[7] The list includes indirect owners through chained ownership. That is, if a firm j is owned by another corporation k, we assign the owners of corporation k as the owners of firm j.

The PSIS records individual-level enrolment and graduation at all public and private not-for-profit post-secondary institutions in Canada. It notably includes information on which institution a student is enrolled in, the program and credential types, as well as the field of study, coded using 6-digit Classification of Instructional Programs (CIP) codes.

Sample Selection The procedure we use to link children to their parents mimics the one Statistics Canada used to create the Intergenerational Income Database (Corak and Heisz, 1999). Using T1 Family Files (T1FFs), we link 15 to 19 year-olds (children) to their parents using unique census family identifiers.[8] Since the CEEDD coverage starts in 2001, the earliest birth cohort for which child-parent linkages are feasible is 1982 (19 year-olds in 2001).

Our main analytical sample comprises all linked children of the 1987-1989 birth cohorts, although we do include older cohorts in some auxiliary analyses. The main reason for focusing on the 1987-1989 cohorts is that graduation data from PSIS only starts in 2009. Many students complete a Bachelor's degree when they are 21 or 22 year old. Then, many children born in 1982 would have graduated many years before the PSIS coverage starts, making it impossible to infer their education. In contrast, coverage of educational outcomes for later birth cohorts is excellent, but we don't observe these cohorts' income at later ages. Our choice of cohorts is meant to balance the trade-off between coverage of education outcomes and coverage of later-life income. Children born in 1987-1989 were between the ages of 20 and 22 in 2009, when PSIS starts, and were between the ages of 29 and 31 in 2018, when the coverage of tax data ends. This main sample includes just over 970,000 children.

Variables Definitions Our definition of income variables largely follows prior work on intergenerational income mobility in Canada (Haeck and Laliberté, 2024; Corak, 2020; Connolly et al., 2019). We define total income as income from all sources before tax (earnings, interest and investment income, self-employment net income, taxable capital gains/losses and dividends, and benefits).[9] When studying cross-sectional differences in income, we measure children income as the average between the ages of 25 and 29 inclusively to reduce the influence of transitory annual income shocks. Parental income is defined as the average total family income (the sum of both parents' income) when a child is between the ages of 15

[8]The exact matching procedure is described in Appendix A.1.

[9]In prior work based on Canadian tax data, the income of individuals who do not file taxes in a given year (i.e. who have no T1 form) is often assumed to be zero (e.g. Connolly et al. (2019)). Using the T4-ROE files, which are issued by employers, we find that there is a non-negligible fraction of employed individuals (roughly 5%) who do not file taxes despite earning taxable income. In those cases, we inpute their income using their unreported T4 earnings.

and 19. We convert average income into percentile ranks for both parents and children, and percentiles are calculated separately for each children birth cohort.

From PSIS, we define education groups as unique combinations of a field of study (4-digit CIP code), a degree type, a post-secondary institution, and an indicator for graduating from (as opposed to not completing) the program. An example of such a group would be graduates of a Bachelor's degree program in civil engineering from the University of Toronto - St George campus. For workers with no post-secondary education, a group is a locality (census subdivision), which proxies the set of high schools they attended. The locality is assigned on the basis of the child's first ever recorded place of residence in the tax data. Overall, children in our main sample belong to just over 14,000 unique education groups.

3 Intergenerational income mobility and employers

3.1 Methods and Measurement

Econometric specification. Following Haeck and Laliberté (2024), one can summarily evaluate the role of employers for the transmission of income across generations (in an accounting sense) using a rank-rank regression of the form

$$y_i = \sum_{p=1}^{100} \beta_p 1\{x_i = p\} + \sum_j \delta_{j(i)} + \varepsilon_i \tag{1}$$

where y_i is the income rank of child i working in firm j, x_i is the income rank of their parents, and $\delta_{j(i)}$ is a complete set of firm fixed effects, which is normalized to have a mean of zero in the estimation sample.

The average income rank of children with parental income p, $\bar{y}_p = E[y_i|x_i = p]$, can be decomposed into a component reflecting within-firm income persistence (β_p), and one capturing sorting across employers by parental income rank (Δ_p):

$$\bar{y}_p = \underbrace{\sum_j \left(\bar{y}_{jp} - \hat{\delta}_j\right) s_{j|p}}_{\beta_p} + \underbrace{\sum_j \hat{\delta}_j s_{j|p}}_{\Delta_p} \tag{2}$$

where $\bar{y}_{jp}$ is the average income rank of children in firm j belonging to income group p, and

$s_{j|p}$ is the share of children of group p who are employed at firm j. This regression-compatible decomposition (Fortin, 2008) neatly separates income gaps into within-firm income gradients and sorting into firms. In the language of typical Kitagawa-Oaxaca-Blinder-type decompositions, the firm sorting effect Δ_p is the "explained" component.[10]

There are three key sets of parameters to study. First, the distribution of employers across groups ($s_{j|p}$) chiefly influences the sorting component.[11] Second, the estimated firm effects $\hat{\delta}_j$ govern the relative weight put on the within- and between-firm components. For example, if compensation was similar across employers (i.e., if the variance of $\hat{\delta}_j$ is close to zero), then firm sorting by parental income would be inconsequential for income gaps and the between-firm component would be small. In section 3.3, we separate the role of firm-specific pay premiums from that of selection on unobserved earnings by plugging in selection-adjusted estimates of $\hat{\delta}_j$. Finally, within-firm differences in average income $\bar{y}_{jp}$ drive within-firm income persistence. We unpack the mechanisms driving differences across parental income groups by examining how education and social connections shape the distribution of firms ($s_{j|p}$), as well as within-firm income differences ($\bar{y}_{jp}$), in sections 4 and 5, respectively.

Assignment of main employers. Since child income is averaged over a five-year window, we take each child's modal employer during that period to assign firms cross-sectionally. Implicitly, this means we match children to their most stable early career job.[12] Similarly, prior work on parental connections generally focus on children's first stable job (Staiger, 2022; San, 2022; Kramarz and Skans, 2014). This procedure assigns children in our estimation sample to over 230,000 different firms.

Several children have no employer throughout the 5-year observation window. Figure A1, panel A shows the fraction of children without an employer for each percentile of the parental income distribution. Close to 20% of children from the bottom 5 percentiles have no

[10]For example, the part of the income gap between children from families at the 25th and 75th percentile of the income distribution that is explained by differences in the firms they work at is $\sum_j \hat{\delta}_j \left(s_{j|75} - s_{j|25}\right)$.

[11]If workers sort into firms on the basis of group-based comparative advantage, then the distribution of firms may also contribute to observed within-firm gaps. For instance, we can write $\beta_p = \sum_j \left(\bar{y}_{j|p} - \delta_j\right) s_j + \sum_j \left(\bar{y}_{j|p} - \delta_j\right)\left(s_{j|p} - s_j\right)$. In our data, the correlation between $(\bar{y}_{j|p} - \delta_j)$ and $(s_{j|p} - s_j)$ is only -0.04, suggesting that, if anything, groups of workers tend to be under-represented at firms at which they experience the largest earnings advantage relative to other workers at that firm, on average. We further discuss this point in Appendix B.2.

[12]Because children can have more than one employer between the ages of 25 and 29, the firm fixed effects we estimate reflect both differences in compensation across firms as well as firms' influence on outside job opportunities. For example, firms that provide good training and networking opportunities may help their workers advance their careers at other employers. Firm fixed effects should therefore be interpreted as reduced-form differences in income between people who had different main employers.

employer. The relationship between non-employment and parental income is not monotonic: while it decreases with parental income over most of the distribution, it increases slightly among the top percentiles. This is mostly due to many children of high-income families receiving business income rather than employment income (see Figure A1, panel B). Hence, we assign children with no employer to either a "non-participant" category or a "capitalist" category, depending on whether they receive any form of capital income (business income, dividends, interest payments).[13] We discuss the implications of the employment margin for intergenerational income mobility in the next section.

Since the estimation dataset is restricted to individuals from only three birth cohorts, there are numerous singleton firms – firms at which we only observe one worker from our sample of children. These firms typically have more than one worker, but only one of them belongs to the 1987-1989 birth cohorts. In those cases, the firm fixed effect cannot be estimated in equation (1). Since these small firms may still play an important role for income mobility, rather than dropping these observations we group them for the purpose of estimating δ_j.[14] The groups are defined as ventiles of the distribution of AKM firm fixed effects estimates. The estimation details of the AKM specification are provided in section 3.3, below. It is worth noting that if we estimate equation (1) dropping singleton firms, the correlation between the estimates of $\hat{\delta}_j$ we obtain and the corresponding AKM firm effects $\hat{\psi}_j$ is strikingly high, at 0.8.

3.2 Main Results

In Figure 1, panel A, we report estimates of the unconditional rank-rank relationship ($\bar{y}_p$), and the estimated average ranks conditional on employers ($\hat{\beta}_p$). The linear slope of the unconditional relationship is 0.232, very similar to prior work using Canadian data (Connolly et al., 2019; Connolly and Haeck, 2024; Haeck and Laliberté, 2024).[15] There are important non-linearities at the tails, with the slope steepening below the 20th percentile as well as at the very top of the parental income distribution. These features are not unique to Canada – very similar non-linearities at the bottom are found in Australia (Deutscher and Mazumder,

[13]Since the decision to participate in the labor market depends on one's labor market opportunities, we keep all children in our sample to avoid sample selection issues.

[14]Appendix Figure A3 shows rank-rank relationships excluding all singleton firms from the sample. The series excluding children at singletons firms lie slightly below those including them, which means children at singleton firms earn a little less, on average. This level shift is apparent throughout the parental income distribution.

[15]Appendix Figure A4 compares the rank-rank relationship in our sample with the rank-rank relationship for different Canadian samples and datasets.

Figure 1: Intergenerational Income Mobility and Children's Employers

Panel A: Child-Parent Income Rank-rank Relationship

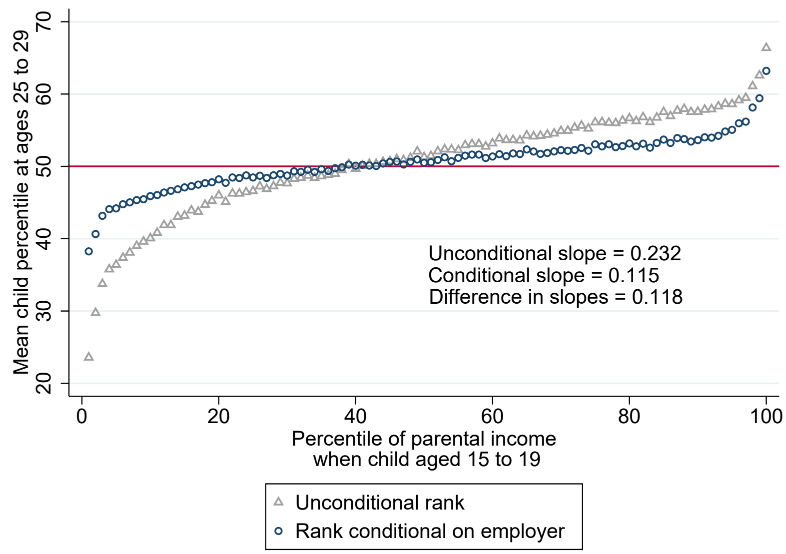

Panel B: Firm Sorting Component

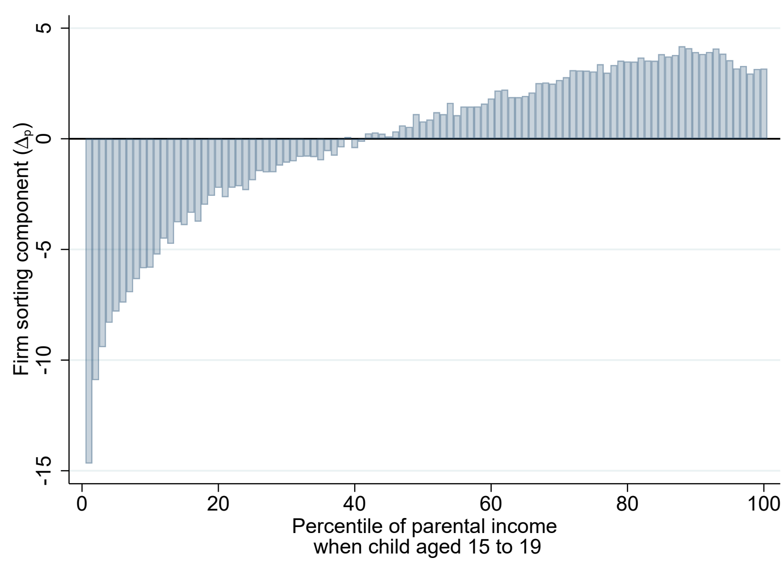

Notes: Panel A shows mean child income percentiles for each parental income decile. Grey triangles show unconditional means, whereas blue circles indicate conditional means accounting for differences in employers, as per equation (1). The difference in linear slopes represents the contribution of employers to income mobility. Panel B shows the firm sorting component. This component is equal to the vertical distance between the two series presented in panel A.

2020), in France (Kenedi and Sirugue, 2023), and in the US to a lesser extent (Chetty et al., 2014). In contrast, the steepening of the rank-rank relationship at the top is even more pronounced in Norway and Sweden (Bratberg et al., 2017).

The linear rank-rank slope conditional on firm fixed effects is 0.115, indicating that the within- and between-firm components are roughly equally important. Interestingly, the non-linearity at the top end of the parental income distribution is mostly a within-firm phenomenon. In contrast, the shape of the conditional rank-rank relationship at the bottom of the parental income distribution differs considerably from that of the unconditional relationship. That is, the steeping of the rank-rank relationship below the 20th percentile is mostly accounted for by firm sorting. To illustrate this with greater visual clarity, in panel B we plot the firm sorting component, which corresponds to the vertical distance between the conditional and unconditional rank-rank relationships, and is equal to the average firm fixed effect $\hat{\delta}_j$ for children of parental income group p. It shows a sharp steepening of firm sorting component below the 20th percentile of the parental income distribution. In contrast, the slope turns negative above the 95th percentile. If anything, children from the richest (top 5 percent) families appear to work at firms that pay less, on average, than the firms children from the 85th to 95th percentiles work at.

When interpreting these results, it is important to keep in mind that the firm sorting component partly reflects differential sorting into employment. Appendix Figure A2 compares the rank-rank relationship in the full sample with the corresponding relationship for the sub-sample of children with an employer. Both series are very similar. Although it is less pronounced than in the full sample, a significant steepening of the firm sorting component below the 20th percentile remains even when conditioning on the employment margin.

Robustness to measurement at older ages. Since several children may still be pursuing graduate degrees between the age of 25 and 29, we might mis-characterize the role of firms for lifetime mobility when measuring income and employers too early. To examine to what extent that might be the case, we reproduce our main decomposition using the 1984-86 birth cohorts, for which we can observe income and employers at older ages, but cannot observe education. For benchmarking purposes, in Figure A5 we first compare the rank-rank relationships for average income between the ages of 25-29 between our two subsamples of birth cohorts. The two series align very closely, although the rank-rank slope is a bit smaller for the 1984-86 cohorts (0.222) than for the 1987-89 cohorts (0.232), consistent with prior work showing that income persistence has been increasing over time in Canada (Connolly and Haeck, 2024).

Figure A6 shows both the unconditional rank-rank relationship as well as the estimated ranks conditional on main employer for ages 25-29, 27-31 and 29-33, using the 1984-86 birth cohorts. While the unconditional relationship steepens as we measure outcomes at older ages, the relative importance of firms remain fairly stable, around 50% in all cases.

Lastly, we examine whether early career main employers are predictive of future income gaps. Perhaps early career jobs are mostly transitory and unrelated with future employment prospect. Alternatively, it might be Naturally, it might also be that for many people their main employer in their early 30s is the same one they had in their late 20s. Here, we use average income between the ages of 29-33 as our outcome, but control for children's main employer at 25-29. The resulting average ranks conditional on employers, shown in Figure X, are very similar to those obtained using the main employer at 29-33. This suggests early career employers are pivotal for one's career.

Excluding non-employment income. Our measure of income includes all sources of income. While firms may contribute to workers' non-labor income – e.g. by offering stock options – the bulk of the compensation associated with an employment relationship is in the form of wages and salaries. That is, employers may play a larger role for mobility in terms of labor income than in terms of total income. We therefore reproduce our main decomposition analysis using employment income (i.e., T4 earnings) only, in Figure A8.[16]

We find that the unconditional rank-rank slope based on children's employment income alone is 0.204, that is 12% lower than the slope based on total income. Strikingly, the rank-rank relationship for employment income turns negative above the 95th percentile of parental income. This means the income advantage of children from very rich families relative to that of families just below them in the distribution is entirely driven by non-labor income. Comparing the linear slopes for the unconditional and conditional relationships, we find that employers account for (0.122/0.204=) 60% of differences in employment income across parental income groups.

[16] For ease of comparability, we do not generate new income percentiles based on the distribution of employment income alone, but rather keep the percentile cutoffs from the original distribution of total income. We then uniformly re-allocate the total amount of non-employment income among children in our sample. That is, we obtain percentile ranks for a counterfactual scenario that holds the total amount of income in our sample fixed, where each child received the same amount of non-labor income and so employment income is the only source of variation in income across children.

3.3 Between-firm component: firm premiums or selection?

Our baseline firm fixed effects $\hat{\delta}_j$ conflate the causal effect of working at firm j (i.e. firm j's value-added) and selection in firm j on the basis of unobserved earnings potential. While the latter is important for inequality accounting, the interpretation of counterfactual analyses hinges on whether estimates of firm effects are causal or not. For instance, if we want to know how the rank-rank slope would change under policies that affect the allocation of children to employers (e.g. policies that reduce the scope for use of social connections in hiring), causal estimates of firm effects are necessary. To this end, we use AKM worker and firm effects to adjust our income rank-based estimates of δ_j for compositional differences in time-invariant unobserved earnings potential.

We estimate the usual AKM two-way (worker and firm) fixed effect model of log earnings on the full sample of 24-54 year-old workers for years 2001-2018.[17] As in prior work, we restrict estimation to the largest connected set (Card et al., 2013). For comparability purposes, we impose additional sample restrictions. First, we drop all observations for which a worker made less than 5,000$ (in 2016 Canadian dollars) in annual earnings.[18] This restriction is meant to drop any observation for which a worker spent very little time in the labor market within a fiscal year (e.g. if they started working in December). An implication of this restriction, however, is to make the AKM firm effect for employers that hire many workers at very low pay scales more positive than it would be otherwise. Second, we drop all workers who appear in the tax data for less than 3 years. This insures that all workers in the sample spent at least one full year in the labor market.

For robustness purposes, we consider two complementary approaches using AKM worker α_i and firm effects ψ_j. The first one casts the problem of obtaining selection-adjusted firm effects in eq. (1) in terms of the the omitted variable bias, where the omitted variable is person i's earnings potential. Assuming AKM worker effects are good proxies for unobserved earnings potential, we replace the parental income dummies in eq. (1) with a cubic polynomial in AKM worker effects.[19] Hence, the firm fixed effects we estimate this way are adjusted for compositional differences across firms in terms of AKM worker effects. Rather

[17] See Appendix B.3 for the econometric specification.

[18] Similarly, Lachowska et al. (2020) drop workers with less than 2,850$ (in 2005 US dollars) in annual earnings, which is roughly the same cutoff as ours for a 1.23 CPI adjustment (translating 2005 US dollars into 2016 US dollars) and at an exchange rate of 1.35.

[19] We also include a dummy variable for individuals who do not have a AKM worker effect. This will be the case for people who never work, or for those who only work for firms outside the largest connected set, for instance. It is worth noting that some people with no main employer between the ages of 25 and 29 may still have a AKM worker effect if they had jobs at other ages.

than grouping singleton firms like we do in our baseline specification, here we drop them and inpute their values in a second step using information on AKM firm effects. More precisely, we predict values of selection-adjusted firm fixed effects for singletons firms by projecting the estimated $\hat{\delta}_j$'s (obtained controlling for AKM worker effects) on AKM firm effects.[20]

Under the second approach, we generate counterfactual income ranks for each child i in our sample by imposing that everyone receives the same AKM firm effect. We then take the difference between their real and counterfactual income ranks – which constitutes the firm's contribution to their income rank – and calculate the average of this difference for each employer.[21] This approach maintains a key property of the AKM log additive specification, which is that high-pay firms increase earnings *in dollars* more for high-pay workers than for low-pay ones. Since AKM firm effects don't exist for the two non-employment categories, as well as for firms outside the largest connected set, we fill these holes using the estimates from the first approach described above.

Figure 2: Robustness to alternative estimates of firm premiums

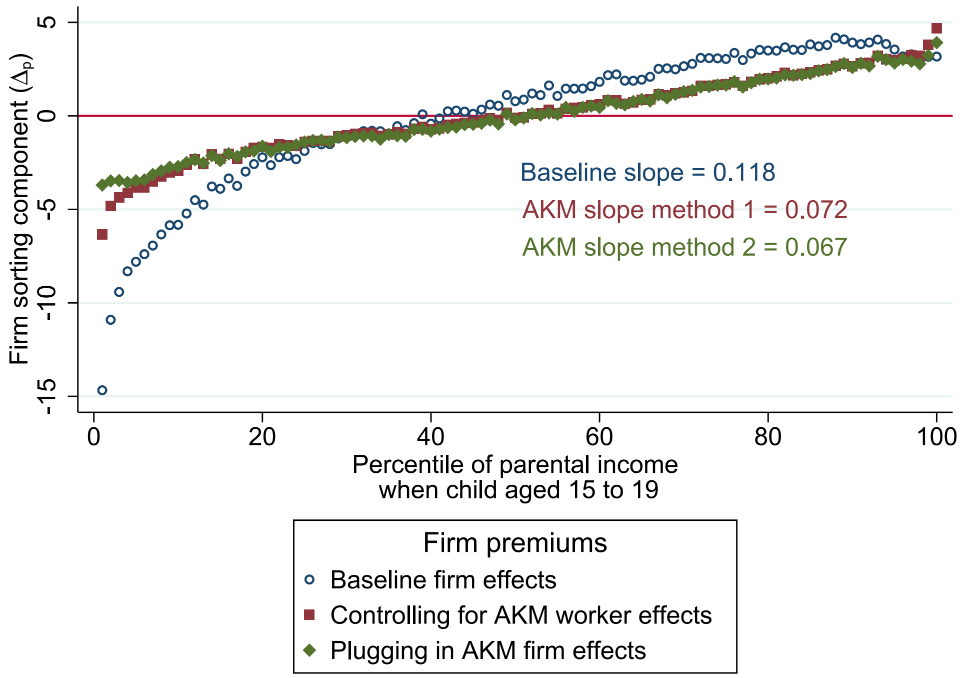

Notes:

Figure 2 shows the firm sorting component Δ_p with selection-adjusted firm effects. Using

[20]The correlation between the selection-adjusted firm effects we obtain under this procedure (excluding singleton firms) and corresponding AKM firm fixed effects $\hat{\psi}_j$ is very high, at 0.95.

[21]Implementation details are provided in Appendix B.4.

the linear rank-rank slope as a summary measure of the firm sorting component, we find the slope decreases from 0.118 to 0.072 when controlling for AKM worker effects and to 0.067 when directly using AKM firm effects. That is, employer-specific pay premiums explain close to a third (e.g. $0.072/0.232 \approx 31\%$) of the transmission of income across generations. In comparison, Dobbin and Zohar (2023) find that AKM firm effects account for 25% of the intergenerational income elasticity in Israel.

Interestingly, not only is the slope of the selection-adjusted series flatter, but the patterns observed at the tails of the parental income distribution are significantly different. First, there is almost no non-linearity at the bottom for selection-adjusted measures. This means the sharp steepening of the rank-rank relationship at the bottom mostly reflects selection effects. We note, however, that the approach based on controlling for AKM worker effects shows slightly more non-linearity at the very bottom, likely because it is less affected by the AKM sample restriction on minimal annual earnings. That is children from the bottom 20 percentiles may be over-represented at low-paying firms that systematically pay some workers less than 5,000$ a year. Second, there appears to be slight patterns of negative selection into firms at the very top of the parental income distribution.

In the next section, we set aside the distinction between selection and causal firm effects and investigate the reasons why children of different background are under- and over-represented at different firms.

4 Quantifying the role of education and social connections for firm sorting

Why are there parental income gaps in access to high-paying employers? We consider two broad types of explanations: differences in human capital and differences in social connections.[22]

[22] Another possible explanation would be differences in preferences for workplace characteristics. Our data does not allow us to credibly examine that scenario.

4.1 Empirical Specification

Our objective is to construct counterfactual distributions of workers across firms $\tilde{s}_{l|p}$ by predicting the probability that a worker from group p matches with firm l, using their education and social connections as predictors.[23]

Our starting point is a dyadic set-up:

$$H_{il} = \alpha_{e(i),l} + \gamma C_{il} + \epsilon_{il} \tag{3}$$

where $e(i)$ denotes the education group worker i belongs to, and H_{il} is an indicator variable for worker i being employed by firm l. $\alpha_{e(i),l}$ is a match effect between education program e and firm l, and C_{il} is a set of variables indicating whether worker i is socially connected to firm l. We consider different types of connections and include them simultaneously as separate predictors. The first one is based on parents' employers, where the connection indicator takes a value of one if at least one of worker i's parents was employed by firm l at any point within the 5-year window we use to find children's modal employer. The second type of connection we consider indicates whether at least one of worker i's parents owns firm l at some point within the 5-year window.

The coefficient γ reflects the impact of parental connections on hiring probabilities. Given the inclusion of education-by-firm fixed effects $\alpha_{e(i),l}$, the coefficient γ is identified by comparing "classmates" (i.e., people who hold the same degree from the same institution) to each other, as in Kramarz and Skans (2014). Put differently, classmates are used to infer a worker's job opportunities in the absence of social connections.

Then, with estimates of $\alpha_{e(i),l}$ and γ in hand, we can calculate counterfactual distributions – predicted shares of children of group p employed at firm l – by aggregating fitted values at the parental income group level:

$$\tilde{s}_{l|p} = E[\hat{\alpha}_{e(i),l}|l,p] + E[\hat{\gamma} C_{il}|l,p] \tag{4}$$

To better understand the intuition behind the predictions obtained via equation (4), consider the special case where social connections are excluded from the model (i.e. $\gamma = 0$). Then, the approach described by equations (3) and (4) is equivalent to a simple reweighting

[23] Whereas $j(i)$ denotes the firm worker i actually works for, here l indexes firms independently of the actual employee-employer match.

procedure in which $\tilde{s}_{j|p} = \sum_j s_{j|e}\nu_{e|p}$, where $s_{j|e}$ is the share of workers with education e working at firm j, and $\nu_{e|p}$ is the share of children with parental income p who have education level e (Haeck and Laliberté, 2024).[24] This simplified exercise informs what the distribution of children across firms would look like if only education mattered – i.e. if classmates all had the same match probabilities, and so differences across parental income groups were only driven by differences in education.

Finally, the associated counterfactual sorting component is

$$\tilde{\Delta}_p = \sum_l \hat{\delta}_l \tilde{s}_{l|p} = \sum_l \hat{\delta}_l E[\hat{\alpha}_{e(i),l}|l,p] + \sum_l \hat{\delta}_l E[\hat{\gamma} C_{il}|l,p]. \tag{5}$$

Note the model in equation (3) does not include information on parental income. That is, we predict individual assignment to firms using information on education and connections alone, and then examine the extent to which the counterfactual sorting component $\tilde{\Delta}_p$ can reproduce the true firm sorting component Δ_p. Equation (5) also makes clear that the quantitative importance of connections for the firm sorting component of income mobility depends on the effects size (i.e. the magnitude of $\hat{\gamma}$), but also on the incidence of connections (the values of C_{il}).

4.2 Estimation

The estimation dataset for the dyadic set-up of eq. (3) is a list of all possible pairwise combinations of workers and firms.[25] With hundreds of thousands of workers and hundreds of thousands of firms, this makes several billions of possible pairwise combinations of workers and firms. As a result, direct estimation of eq. (3) on the full dataset is computationally impossible.

Fortunately, it remains possible to retrieve all parameters from eq. (3). An important observation is that education-firm pairs with $\bar{C}_{e,l} = 0$ (i.e., no one from education group e is connected to firm l) do not contribute to the identification of γ (Kramarz and Skans, 2014).[26] We can therefore estimate γ using only the subsample of education-firm pairs for which $\bar{C}_{e,l} >$

[24] One can always decompose the true distribution $s_{j|p} = \sum_j s_{j|e,p}\nu_{e|p}$. The reweighting procedure amounts to replacing $s_{j|e,p}$ with $s_{j|e}$ in the summation.

[25] The set of firms is restricted to those that hire at least one child in our sample. We therefore do not predict probabilities of matching with out-of-sample firms that do not hire anyone from our sample of children.

[26] The "no employer" groups fall into this category since no one is connected to the absence of an employer.

0 without loss of identifying variation. That is, estimates of γ for the full sample are exactly the same as those one obtains from estimation on the subsample of education-firm pairs for which $\bar{C}_{e,l} > 0$. Still, that estimation subsample remains quite large. For computational reasons, we therefore estimate eq. (3) separately by broad categories of education. We combine the 14,000 education groups e into just over 300 categories (indexed by m) based on broad fields of study (2-digit CIP codes), and estimate a value of γ^m separately for each category.[27]

With estimates of γ^m in hand, we can then back-out the values of $\alpha_{e(i),l}$ for all education-firm pairs (including those not in the $\bar{C}_{e,l} > 0$ subsample), calculate the corresponding fitted values based on eq. (3) for all possible worker-firm pairs, and average them by parental income group to obtain predicted distributions using eq. (4). The method for retrieving the fitted values is described in Appendix B.1.

4.3 Estimates of Social Connection Effects

Main estimates. Figure 3 plots the distribution of estimates of the parental connections coefficients γ^m, separately for connections based on a parent working at a firm and for connections based on a parent owning a firm. Most of the estimates of the change in the probability of working for a given employer if one's parent works there range between 0.024 (5th percentile) and 0.064 (95th percentile), with a mean of 0.041. For comparison, the main estimate of that same parameter is 0.061 in Kramarz and Skans (2014), who use Swedish data. Our results also indicate that the impact of parental connections based on ownership status is an order of magnitude larger than for connections based on employment relationships. The two distributions barely overlap. When a parent owns a firm, the change in their child's hiring probability range between 0.204 (5th percentile) and 0.717 (95th percentile), with a mean of 0.492.

Given that high-income parents are much more likely to own a business than low-income ones, this suggests a potentially important role for ownership types of connections for intergenerational mobility. These findings also have important implications for the measurement of social connections in studies of economic inequality. Business owners often pay themselves a salary, which means they appear to be employed by the firms they own. Without ownership data, analysts may therefore conflate the two types of connections we study, despite

[27] For workers without any post-secondary education, we form categories by combining census subdivisions into census divisions.

Figure 3: Distributions of Estimates of Social Connections Effects on Hiring Probabilities

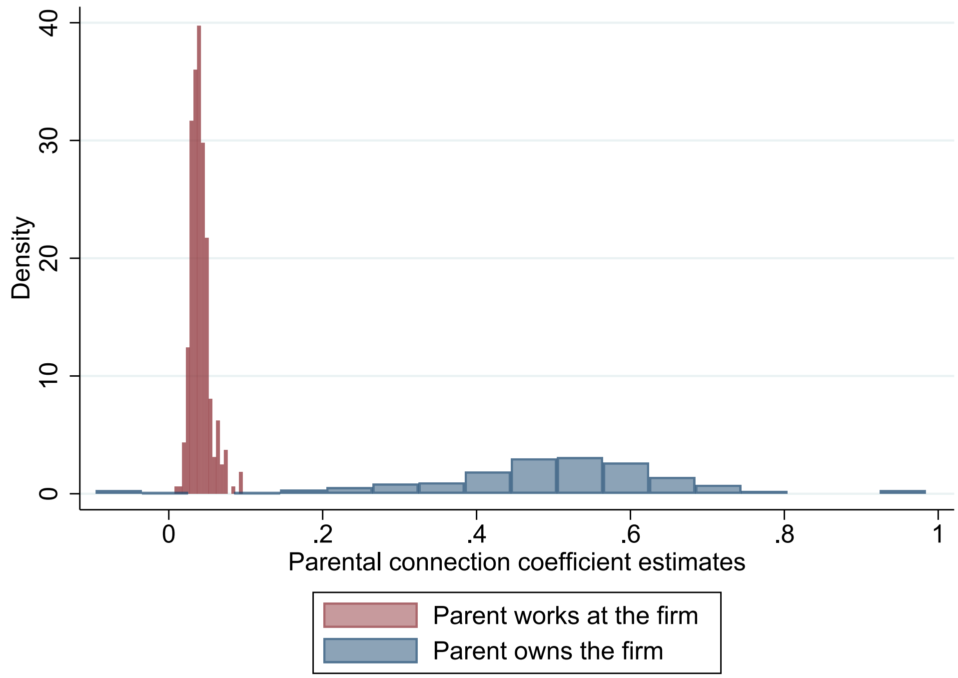

Notes:

their impact on hiring patterns being dramatically different.

In addition, Appendix Figure A9 shows that connections appear to matter more for children without post-secondary education, a result also in line with the findings in Kramarz and Skans (2014). Because children of high-income families are more likely to have post-secondary education, it might be that connections, in particular those based on a parent's employment relationships, are quantitatively more important for low-income families.

Robustness. If parents and children share unobserved firm-specific skills, then the coefficient γ may reflect hiring on the basis of unobserved skills (or unobserved match effects) rather than connections. To verify whether this interpretation is plausible, we re-estimate the dyadic set-up of eq. (3) measuring H_{il} at a specific point in time (when children are aged 25). We then include a set of social connections indicators for different values of the parent's job tenure (at the time the child was 25), including negative tenure, in the estimating equation. That is, we include future connections (i.e. the parent has yet to start working at firm l) as placebo connections. To the extent that worker-firm match effects and firm-specific unobserved skills are time invariant, then the effect of future connections will partly capture hiring by skill.

We estimate the model separately by quantiles of the distribution of $\hat{\gamma}^m$ presented in the previous section. Results are presented in Figure 4. For visual clarity, we only plot coefficient estimates for the top and bottom quantiles of the distribution of $\hat{\gamma}^m$. Positive values of job tenure T mean that a parent of child i worked at firm l when the child was 25, and that they had been employed there for T years at that time. Negative values of job tenures mean that a parent of child i started working at firm l when the child was older than 25. For instance, a job tenure of -2 means the parent started working at the firm when the child was 27.[28]

Figure 4: Social Connections Effects on Hiring Probabilities, Current and Future Connections

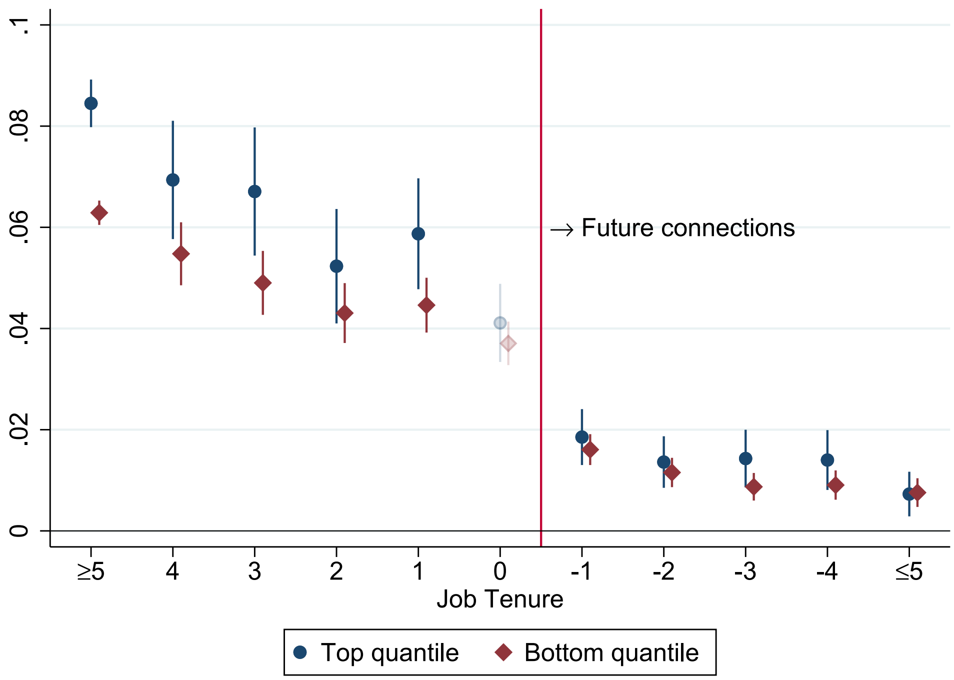

Notes:

The parental connection coefficients are larger the longer the parent has been at the firm. The is a sharp drop in the magnitude of the estimated coefficients between tenures of 1 and -1 years.[29] While the coefficients on future connections do seems to converge toward zero after many years, they tend to remain positive and statistically different from zero at conventional levels. In a similar analysis, Staiger (2022) also find that coefficients on future connections

[28]For this analysis, we do not consider connections based on ownership since a child cannot possibly work at a business that a parent as yet to start (i.e. the firm does not yet exist).

[29]For tenure of 0 year, it is unclear whether the parent joined a firm before of after the child joined their modal employer during the fiscal year.

do not go to zero. It may be the case that the child is the person providing a referral for their parents, or that future parental connections reflect broader contemporaneous social networks – e.g. a connected extended family member who helps both a child and their parent to obtain a job at their firm. However, our results lend little support to that interpretation. If future connections mostly capture broader social networks, then we might expect the coefficients on future connections to be larger in groups where connections matter the most. This is not what we find. The coefficients on future connections are extremely similar for both education categories that have large (top quantile) and small (bottom quantile) estimated $\hat{\gamma}^m$'s. This suggests that while most of the relationship between parental employment and children's hiring probabilities does reflect a social connection effect, we cannot rule out that it may also partly capture some hiring on correlated skills.

4.4 Counterfactual firm sorting components

Next, we examine counterfactual firm sorting components $\tilde{\Delta}_p$. We present these results in Figure 5 by stacking all of the terms in $\tilde{\Delta}_p$. We also tack on the residual term $r_p = \Delta_p - \tilde{\Delta}_p$ to examine how education and connections collectively account for overall patterns of firm sorting.

To quantify the contribution of each dimension to firm sorting, we calculate their share of the area under the curve traced by Δ_p, as in Haeck and Laliberté (2024).[30] This allows to take into account any non-linearity in their explanatory power across the parental income distribution.

We find that the role of education largely exceed that of social connections.

Since neither of the two mechanisms we study have a causal interpretation, some caveats are in order. First, it might be that the role of education partly operates via the use of social networks built in school, rather than to differences in human capital.[31] In that case, part of what we attribute to education may rather reflect a social connection effect.

[30] At some parental income percentiles, the explanatory power of a given factor can be negative. For example, it could be that the sign of Δ_p and that of the education component $\sum_l \hat{\delta}_l E[\hat{\alpha}_{e(i),l}|l,p]$ are opposite. Whenever that arises, we subtract rather than add the area covered by that factor in the calculation of its total contribution to firm sorting across the parental income distribution. This insures that all shares sum to 1.

[31] Zimmerman (2019) finds evidence in support of that mechanism for elite colleges in Chile. Another possibility is that connections are used to gain admission to specific programs or post-secondary institutions, although this is unlikely to happen in Canada since admission to university is only based on grades, with very few exceptions.

Figure 5: Decomposition of firm sorting component

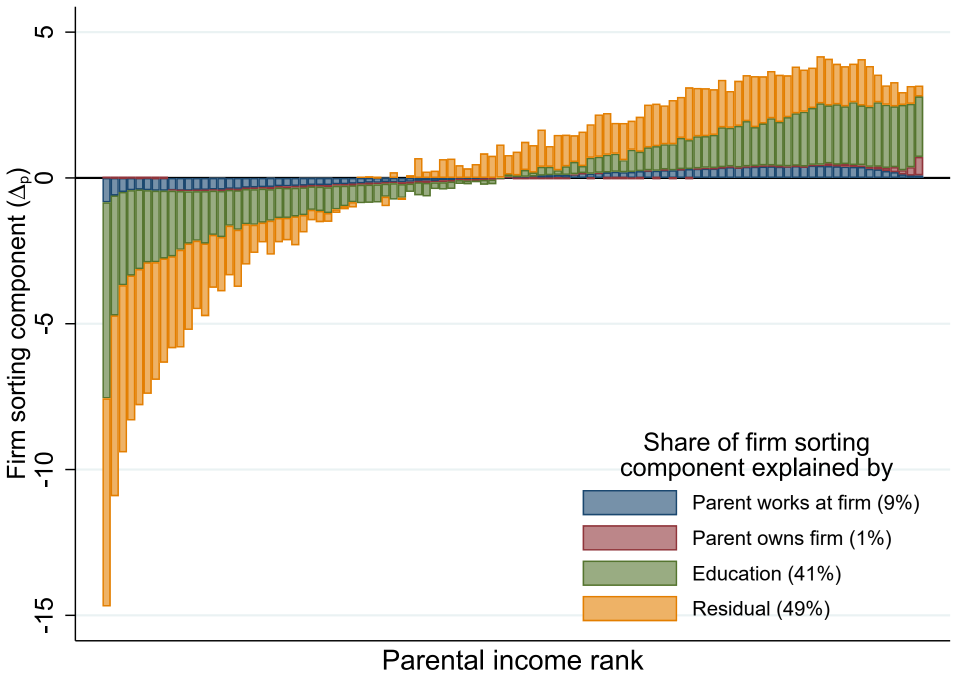

Notes:

However, Eliason et al. (2023) find that connections from former classmates have much smaller effects on hiring than parental connections do. Also, if we omit parental connections from the prediction model altogether (i.e. the role of education is not conditional on parental connections), the explanatory power of education only goes from 41% to 43%. These two pieces of evidence suggest the social-connection aspect of education is likely small.

Second, the residual component is sizable and partly reflects the impact of both unobserved skills and unobserved social connections. That is, we may understate the role of both skills and connections, and it is not clear for which of the two the bias is more severe. Nevertheless, given prior evidence on which types of connections matter most (Eliason et al., 2023), we suspect that even if we were able to fully account for both observable and unobservable dimensions of human capital and connections, the role of education would remain far more important. This is because adding less-important types of connections as predictors is unlikely to increase the explanatory power of social connections by much, whereas the variability in skills among graduates of a given education program is plausibly large.

Nevertheless, social connections could matter greatly for within-firm income gaps, in particular if connected workers receive preferential treatment from their employers.

5 Unpacking within-firm parental income gaps

Why do children of high-income family earn more, on average, than their colleagues from lower income families? To better understand the quantitative importance of different mechanisms contributing to within-firm income differences by parental income, we estimate

$$y_i = \theta C_{ij(i)} + \omega_{e(i)} + \kappa_{j(i)} + \varepsilon_i \tag{6}$$

where y_i is worker i's income rank, C_{ij} indicates whether worker i is socially connected with the firm $j(i)$ that hires them, and $\omega_{e(i)}$ is a set of education group fixed effects. To insure that parameters θ and $\omega_{e(i)}$ are estimated using within-firm variation, we also include firm fixed effects $\kappa_{j(i)}$. However, we also present some results without firm fixed effects to guide the interpretation of findings.

5.1 Income advantages of connected workers

We first report estimates of the social connection coefficients θ in Figure 6. In the top panel, we report estimates from a model without firm fixed effects, in which case the coefficients capture both within-firm income advantages as well as income gains arising from access to better paying firms. While workers who work *with* their parents see a large income advantage of 8 percentile ranks, those who work *for* their parents see a much more modest advantage of 2 percentile ranks. The inclusion of firm fixed effects completely reverses this pattern: children working with their parents don't earn a lot more than their co-workers, but children working at their parents' business have much higher income than their colleagues. This means that the income advantage associated with working with one's parents operates mostly through access to better-paying firms, a conclusion that Staiger (2022) also reaches. In the case of working at a firm parents own, the fact that the coefficient with firm effects is larger than the one without firm effects indicates that firms that hire the owner's children are relatively low-pay firms.

These findings raise the question: why do such low-pay firms offer much higher remuneration to the owner's children relative to other same-aged employees?[32] Since education

[32]Because income ranks are based on total income from all sources, it is possible that the income of the owner's children comes in different forms than for other unconnected employees. For instance, children of the owner may also be given shares in the company and derive business income from it. Other employees

Figure 6: Social Connections and Income Ranks

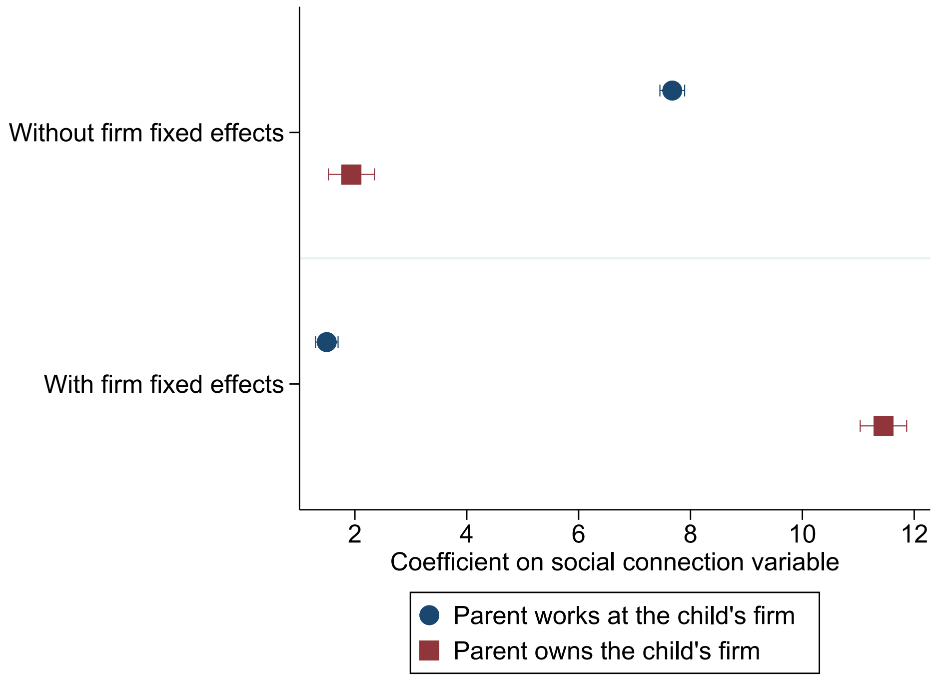

Notes:

is controlled for in the model, the two most likely explanations are that the children either (i) possess a set of unobserved skills that is a particularly good fit for their parent's firm, or (ii) they may be subject to preferential treatment. Regarding the first possible reason, it might be that the child has observed their parent operate the business for many years and so they have accumulated excellent knowledge of the industry and the firm itself. Without individual-level measures of productivity, however, distinguishing between such firm-specific human capital and preferential treatment is a particularly challenging task. To make progress on that question, we flip the perspective and examine whether firms gain from hiring connected workers.

5.2 Parental connections and firm performance

In this subsection, we produce event-study estimates of changes in firm performance around the time of hiring a connected worker. The outcomes we track are log revenue per worker, log value-added per worker, and log payroll per worker. For internal consistency, we

likely only receive employment income.

focus on hiring events when a worker from our main sample of children born in 1987-89 is joining their modal employer (i.e. the firm that they are assigned to in eq. (1)).

Since any individual worker is unlikely to have a detectable impact on performance in large firms, we restrict the sample to private incorporated firms with less than 100 employees. Then, for each firm hiring a connected worker, we find a set of control firms that hired a unconnected worker from the same education group in the same year. That is, both the "treatment" and "control" firms within a comparison set g hired graduates of the same education program e in the same year h, but in one case the worker had a social connection with the employer, and in the other case they did not. Finally, we only keep firms that are observed at least 3 years before and after the hiring event.

The estimating equation is:

$$q_{jt} = \sum_{\tau} \beta_{\tau} \left(\mathbb{1}\{t - h_{q(j)} = \tau\} \times C_j \right) + \phi_j + \phi_{g(j),t} + \epsilon_{jt} \tag{7}$$

where q_{jt} is an outcome for firm j in calendar year t, and C_j indicates whether firm j hired a connected worker (as opposed to an unconnected one). $g(j)$ denotes the comparison set firm j belongs to, and $h_{q(j)}$ is the year all firms in comparison set q hired at least one worker from education group e. ϕ_j are firm fixed effects, and $\phi_{g(j),t}$ are comparison set-by-year fixed effects. The latter account for any time trend that is shared by firms that hired worker from the same program in the same year. We estimate the model separately for parental connections based on employment and ownership relationships.

The coefficients of interest are β_{τ}, which are on interactions between dummies for time relative to the event year h (omitting the dummy for $\tau = -1$), and treatment status C_j. We include consider an event window of $\tau \in [-5, 5]$, but also include observations outside this window and bin the end points (i.e. combining all observations for $\tau \leq -5$ into a single relative-time category).

Results are presented in Figure 7. In the case of firms hiring the child of one of their employees, we detect no significant change in firm performance. Most point estimates are not statistically significant at conventional levels, and no clear trend emerges. The results for firms hiring a child of the owner are completely different. There is no apparent pre-trend, and all performance indices drop on impact. Both revenue per worker and value-added per worker drop by 0.1 log points following the hiring of the owner's child. Payroll per worker decreases too, but less so than revenue. This implies that the revenue to payroll ratio goes down.

The negative impact on payroll per worker does not necessarily contradict prior findings that children of owners have a higher income rank than their similar-aged co-workers. The child that joins their parent's firm may earn less than other (possibly older) co-workers that are not born in 1987-89. It might also be the case that most of the connected child's income is derived in the form of business income rather than employment income. Another possibility is that incumbent workers at the firm may now work fewer hours, and the new employee merely makes up for these lost work hours, so that total payroll is mostly unchanged but the number of workers rises. We will explore these alternative explanations in a future draft.

Overall, the firm responses documented in Figure 7 suggest children of business owners may receive preferential treatment when they work for their parents. That is, the firm owner appears to be willing to incur an economic cost to hire their child, as opposed to an unconnected worker. In the next section, we quantify the importance of the income advantages that accrue to connected workers for intergenerational income mobility.

Figure 7: Firm Performance and Hiring of Connected Workers

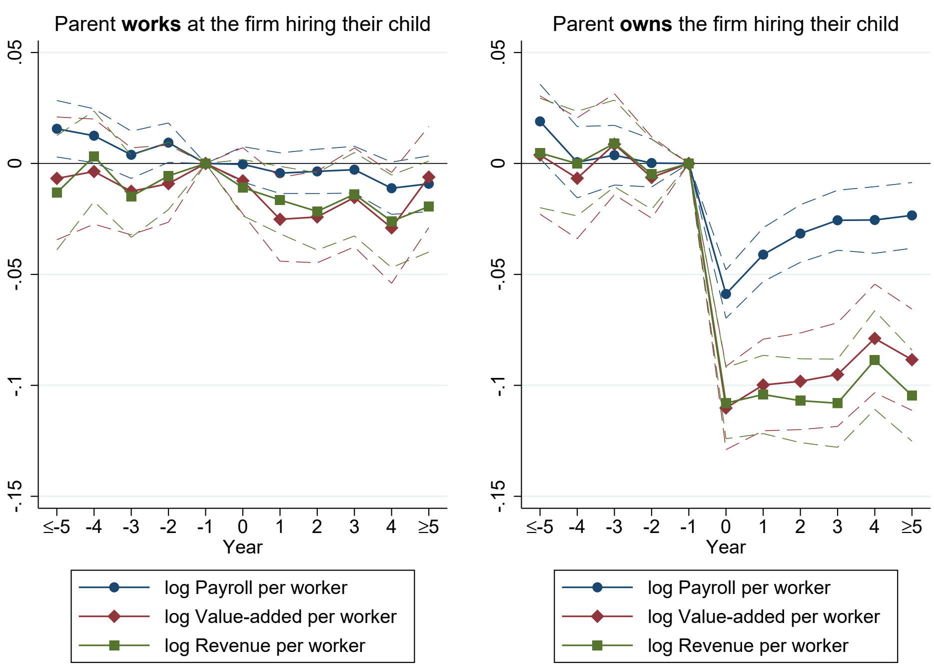

Notes:

5.3 Quantifying factors contributing to within-firm income gaps

Subtracting $\delta_{j(i)}$ on both sides of eq. (6) and averaging by parental income level returns the within-firm component

$$\begin{aligned}\beta_p &= E[y_i - \delta_{j(i)}|p] \\ &= E[\theta C_{ij(i)}|p] + E[\omega_{e(i)}|p] + \big(E[\varepsilon_i|p] + E[\kappa_{j(i)} - \delta_{j(i)}|p]\big)\end{aligned}$$

where $E[\omega_{e(i)}|j,p]$ captures differences in pay explained by differences in education and $E[\theta C_{ij(i)}|j,p]$ captures any income advantage associated with social connections. We combine the last three terms into a residual compensation gap.

Figure 8: Decomposition of within-firm income gaps

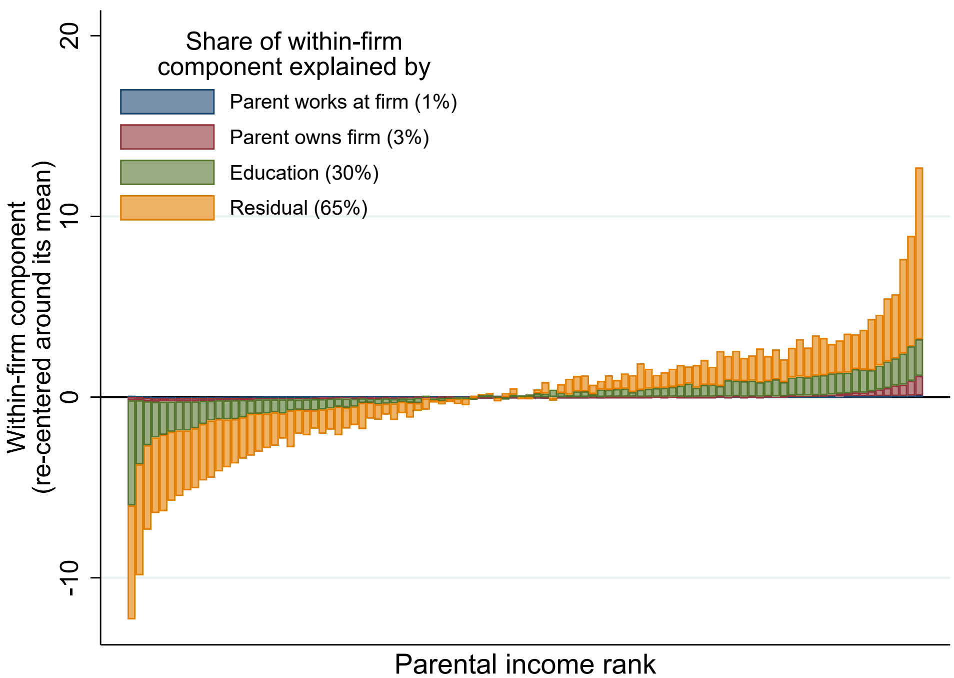

Notes:

Figure 8 presents the results. The contribution of parental connections to employers is minimal, except at the very top where the explanatory power of ownership-based connections comes close to that of education. Overall, differences in social connections (combining both types) only explain 4% of within-firm income gaps.[33] In contrast, differences in education

[33]As in section 4.4, the share explained by a given factor corresponds to the fraction of the total area under the curve.

explain 30% of the rank-rank relationship conditional on employers. Still, most of the variation (65%) is neither explained (in an accounting sense) by connections nor by education. The considerable share of unexplained variation we find is consistent with prior work that shows that most of the variance in AKM worker effects is unrelated to observable measures of productivity.[34] For example, Håkanson et al. (2021) find that the R^2 of a regression of AKM worker effects on cognitive and non-cognitive skills is only 0.24.

6 Discussion and Conclusion

TBA.

[34]Dobbin and Zohar (2023) show that the under the canonical AKM model, the within-firm component of the IGE is entirely due to differences in worker effects.

References

Abowd, John M, Francis Kramarz, and David N Margolis, "High Wage Workers High Wage Firms," *Econometrica*, 1999, *67*, 251–333.

Becker, Gary S. and Nigel Tomes, "An equilibrium theory of the distribution of income and intergenerational mobility,," *Journal of Political Economy*, 1979, *87*, 1153–1189.

Björklund, Anders, Mikael Lindahl, and Erik Plug, "The origins of intergenerational associations: Lessons from Swedish adoption data," *The Quarterly Journal of Economics*, 2006, *121* (3), 999–1028.

Bolte, Lukas, Nicole Immorlica, and Matthew O Jackson, "The role of referrals in immobility, inequality, and inefficiency in labor markets," *arXiv preprint arXiv:2012.15753*, 2020.

Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel D Schnitzlein, and Kjell Vaage, "A comparison of intergenerational mobility curves in Germany, Norway, Sweden, and the US," *The Scandinavian Journal of Economics*, 2017, *119* (1), 72–101.

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The value of hiring through employee referrals," *The Quarterly Journal of Economics*, 2015, *130* (2), 805–839.

Caldwell, Sydnee and Nikolaj Harmon, "Outside options, bargaining, and wages: Evidence from coworker networks," *Unpublished manuscript, Univ. Copenhagen*, 2019, pp. 203–207.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women," *The Quarterly journal of economics*, 2016, *131* (2), 633–686.

— , — , Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

— , Jesse Rothstein, and Moises Yi, "Industry wage differentials: A firm-based approach," *Journal of Labor Economics*, 2024, *42* (S1), S11–S59.

— , Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *NBER Working Paper*, 2012, (w18522).

— , — , and — , "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly journal of economics*, 2013, *128* (3), 967–1015.

Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan, "Income segregation and intergenerational mobility across colleges in the United States," *The Quarterly Journal of Economics*, 2020, *135* (3), 1567–1633.

— , Matthew O Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob et al., "Social capital II: determinants of economic connectedness," *Nature*, 2022, *608* (7921), 122–134.

— , Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where is the land of opportunity? The geography of intergenerational mobility in the United States," *The Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

Chyn, Eric and Lawrence F Katz, "Neighborhoods matter: Assessing the evidence for place effects," *Journal of Economic Perspectives*, 2021, *35* (4), 197–222.

Connolly, Marie and Catherine Haeck, "Intergenerational income mobility trends in Canada," *Canadian Journal of Economics/Revue canadienne d'économique*, 2024.

— , — , and Jean-William Laliberté, "Parental Education and the Rising Transmission of Income between Generations," in "Measuring Distribution and Mobility of Income and Wealth," University of Chicago Press, 2022, pp. 287–316.

— , Miles Corak, and Catherine Haeck, "Intergenerational Mobility between and within Canada and the United States," *Journal of Labor Economics*, 2019, *37* (S2), S595–S641.

Corak, Miles, "The Canadian geography of intergenerational income mobility," *The Economic Journal*, 2020, *130* (631), 2134–2174.

— and Andrew Heisz, "The intergenerational earnings and income mobility of Canadian men: Evidence from longitudinal income tax data," *Journal of Human Resources*, 1999, pp. 504–533.

— and Patrizio Piraino, "The intergenerational transmission of employers," *Journal of Labor Economics*, 2011, *29* (1), 37–68.

Deutscher, Nathan and Bhashkar Mazumder, "Intergenerational mobility across Australia and the stability of regional estimates," *Labour Economics*, 2020, *66*, 101861.

Dobbin, Caue and Tom Zohar, "Quantifying the Role of Firms in Intergenerational Mobility," 2023.

Doepke, Matthias, Giuseppe Sorrenti, and Fabrizio Zilibotti, "The economics of parenting," *Annual Review of Economics*, 2019, *11*, 55–84.

Dostie, Benoit, Jiang Li, David Card, and Daniel Parent, "Employer policies and the immigrant–native earnings gap," *Journal of Econometrics*, 2023, *233* (2), 544–567.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker, "Referral-based job search networks," *The Review of Economic Studies*, 2016, *83* (2), 514–546.

Eliason, Marcus, Lena Hensvik, Francis Kramarz, and Oskar Nordström Skans, "Social connections and the sorting of workers to firms," *Journal of Econometrics*, 2023, *233* (2), 468–506.

Engbom, Niklas and Christian Moser, "Returns to education through access to higher-paying firms: Evidence from US matched employer-employee data," *American Economic Review*, 2017, *107* (5), 374–378.

Engzell, Per and Nathan Wilmers, "Firms and the Intergenerational Transmission of Labor Market Advantage," 2022.

Fortin, Nicole M, "The gender wage gap among young adults in the united states the importance of money versus people," *Journal of Human Resources*, 2008, *43* (4), 884–918.

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative matching or exclusionary hiring? The impact of employment and pay policies on racial wage differences in Brazil," *American Economic Review*, 2021, *111* (10), 3418–3457.

Glitz, Albrecht, "Coworker networks in the labour market," *Labour Economics*, 2017, *44*, 218–230.

Haeck, Catherine and Jean-William Laliberté, "Careers and Intergenerational Income Mobility," 2024.

Håkanson, Christina, Erik Lindqvist, and Jonas Vlachos, "Firms and skills: the evolution of worker sorting," *Journal of Human Resources*, 2021, *56* (2), 512–538.

Hellerstein, Judith K and David Neumark, "Workplace segregation in the United States: Race, ethnicity, and skill," *The review of economics and statistics*, 2008, *90* (3), 459–477.

Kaila, Martti, Emily Nix, and Krista Riukula, "The Impact of an Early Career Shock on Intergenerational Mobility," *Minneapolis Federal Reserve OIGI Working Paper*, 2022.

Kenedi, Gustave and Louis Sirugue, "Intergenerational income mobility in France: A comparative and geographic analysis," *Journal of Public Economics*, 2023, *226*, 104974.

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

Lachowska, Marta, Alexandre Mas, and Stephen A Woodbury, "Sources of displaced workers' long-term earnings losses," *American Economic Review*, 2020, *110* (10), 3231–3266.

Oyer, Paul and Scott Schaefer, "Firm/employee matching: An industry study of US lawyers," *ILR Review*, 2016, *69* (2), 378–404.

Rivera, Lauren A, "Ivies, extracurriculars, and exclusion: Elite employers' use of educational credentials," *Research in Social Stratification and Mobility*, 2011, *29* (1), 71–90.

Sacerdote, Bruce, "How large are the effects from changes in family environment? A study of Korean American adoptees," *The Quarterly Journal of Economics*, 2007, *122* (1), 119–157.

San, Shmuel, "Who Works Where and Why? The Role of Social Connections in the Labor Market," 2022.

Solon, Gary, "Intergenerational mobility in the labor market," in "Handbook of labor economics," Vol. 3, Elsevier, 1999, pp. 1761–1800.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter, "Firming up inequality," *The Quarterly Journal of Economics*, 2019, *134* (1), 1–50.

Staiger, Matthew, "The intergenerational transmission of employers and the earnings of young workers," *Washington Center for Equitable Growth Working Paper*, 2022.

Stuhler, Jan, "A Review of Intergenerational Mobility and its Drivers," *JRC Research Reports*, 2018, (JRC112247).

Weinstein, Russell, "Employer screening costs, recruiting strategies, and labor market outcomes: An equilibrium analysis of on-campus recruiting," *Labour Economics*, 2018, *55*, 282–299.

—, "Firm decisions and variation across universities in access to high-wage jobs: Evidence from employer recruiting," *Journal of Labor Economics*, 2022, *40* (1), 1–46.

Zimmerman, Seth D, "Elite colleges and upward mobility to top jobs and top incomes," *American Economic Review*, 2019, *109* (1), 1–47.

Online Appendix

"Parental Income in the Labor Market"

Jean-William Laliberte and Alexander Whalley

Appendix Figures

Figure A1: Fraction of individuals with no modal employer at age 25-29

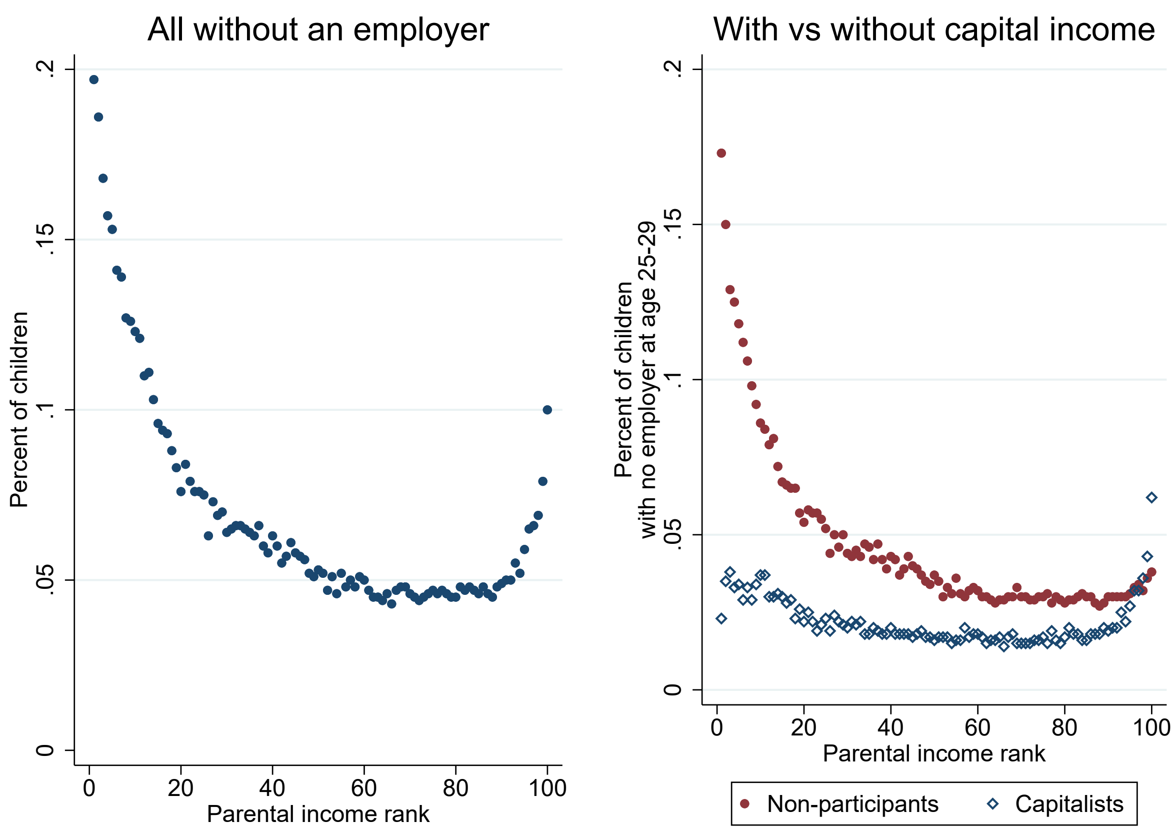

Notes: This figure shows

Figure A2: Income rank-rank relationship, with and without non-employed individuals

Panel A: Unconditional rank-rank relationship

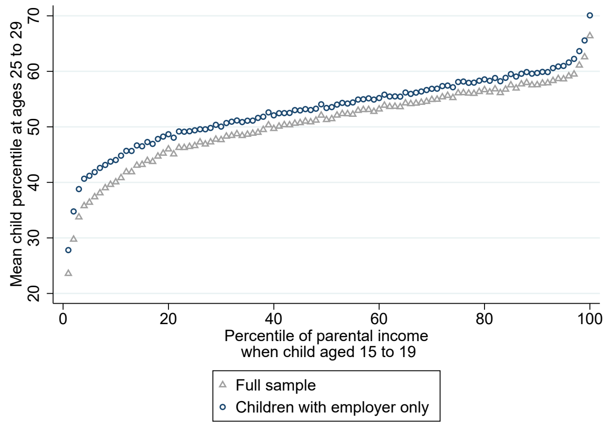

Panel B: Rank-rank relationship conditional on employers

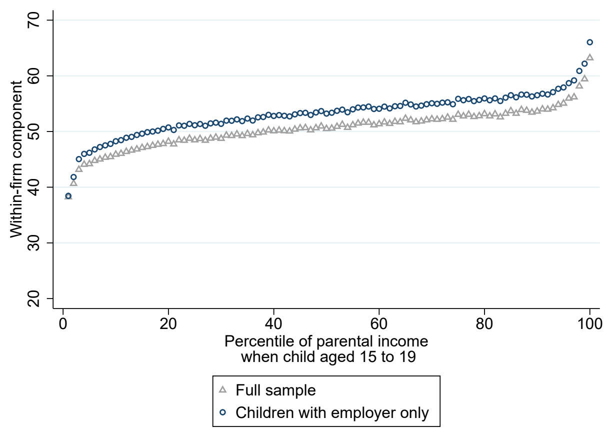

Panel C: Firm sorting component

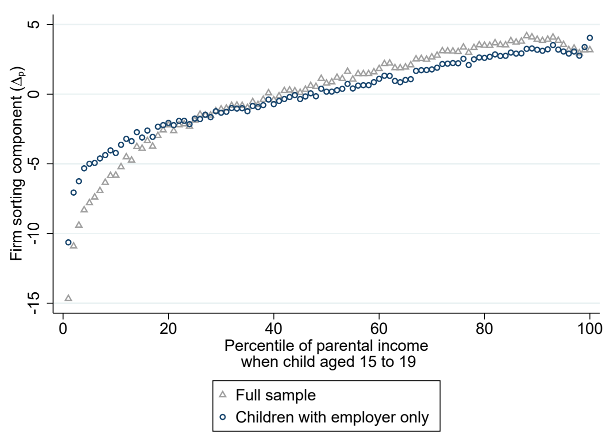

Notes:

Figure A3: Income rank-rank relationship, with and without children at singleton firms

Panel A: Unconditional rank-rank relationship

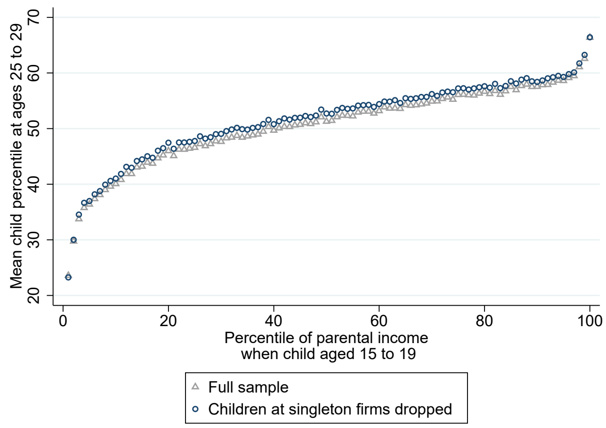

Panel B: Rank-rank relationship conditional on employers

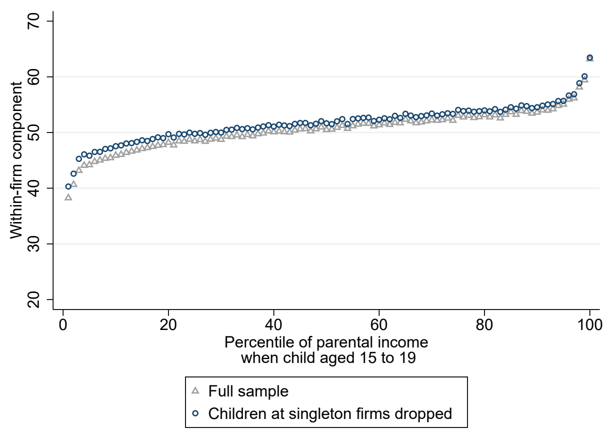

Panel C: Firm sorting component

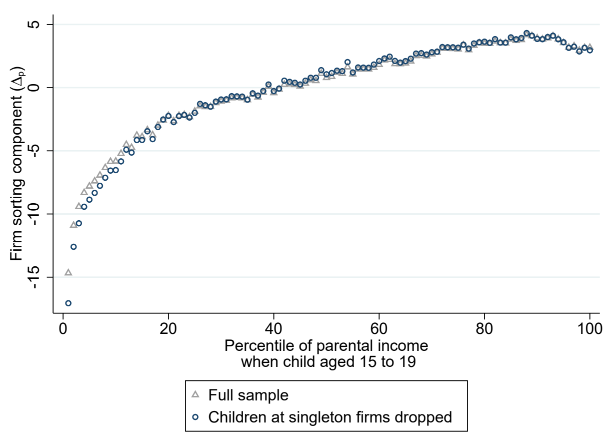

Notes:

Figure A4: Comparison of rank-rank relationships across Canadian datasets

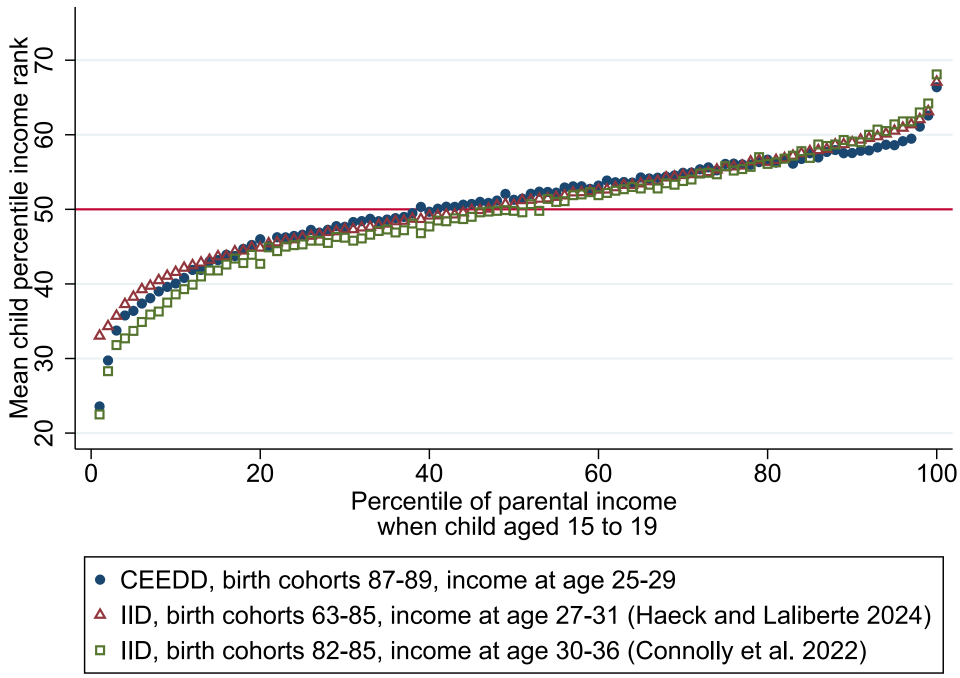

Notes: This figure shows mean child percentile income rank for three different sample. The series in blue circles corresponds to our main sample, drawn from the CEEDD. It includes children born in 1987-89, and their income is calculated when they are aged 25-29. The other two series are based on data from the Intergenerational Income Database (IID). The series in red triangles is from Haeck and Laliberté (2024). It includes birth cohorts from 1963-1985, and their income is calculated when they are aged 27-31. The series in green squares is from Connolly et al. (2022). It includes birth cohorts from 1962-1985, and their income is calculated when they are aged 30-36.

Figure A5: Comparison of rank-rank relationships across birth cohorts

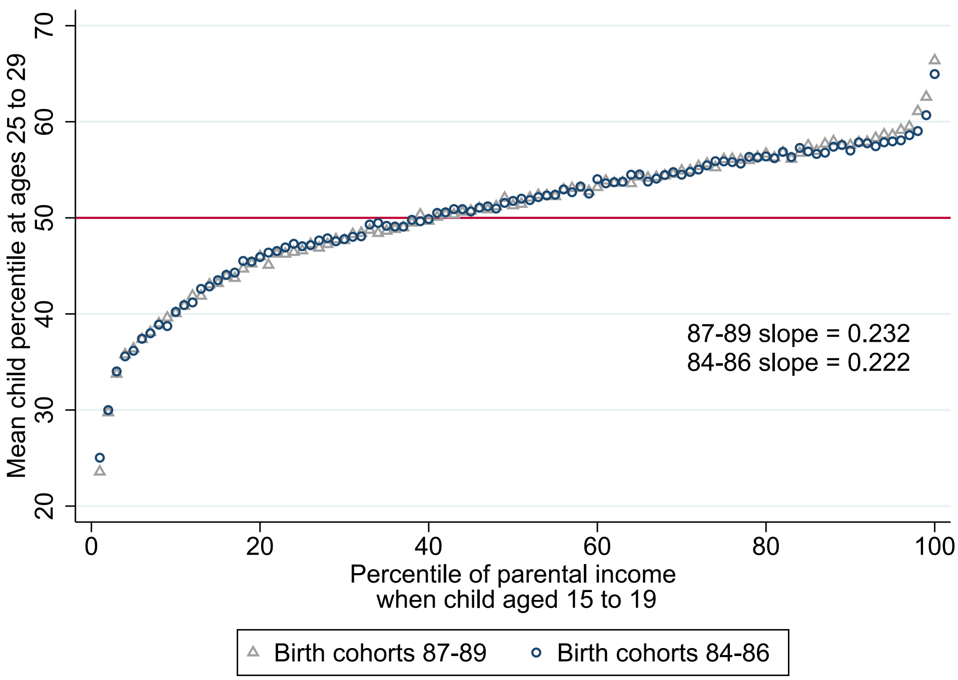

Notes: This figure shows mean child percentile income rank at age 25-29, separately for the 1987-89 and 1984-86 birth cohorts.

Figure A6: Income rank-rank relationship, by age-at-measurement

Panel A: Average income and main employer at 25-29

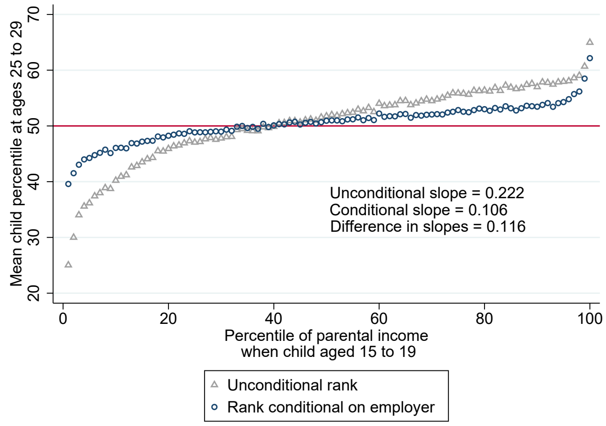

Panel B: Average income and main employer at 27-31

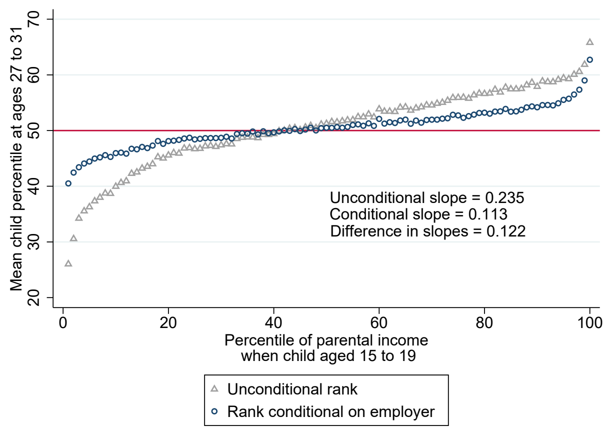

Panel C: Average income and main employer at 29-33

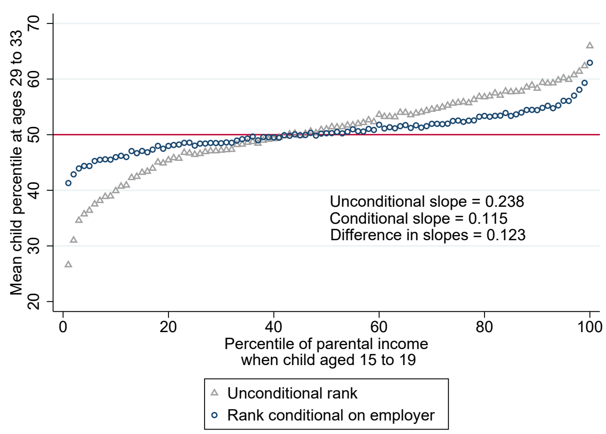

Notes: This figure shows mean child percentile income rank for the 1984-86 birth cohorts. In panel A the outcome is average income rank based on average income between the ages of 25 and 29. Panels B and C reproduce similar analyses using average income at 27-31 and 29-33, respectively.

Figure A7: Income mobility and past employers

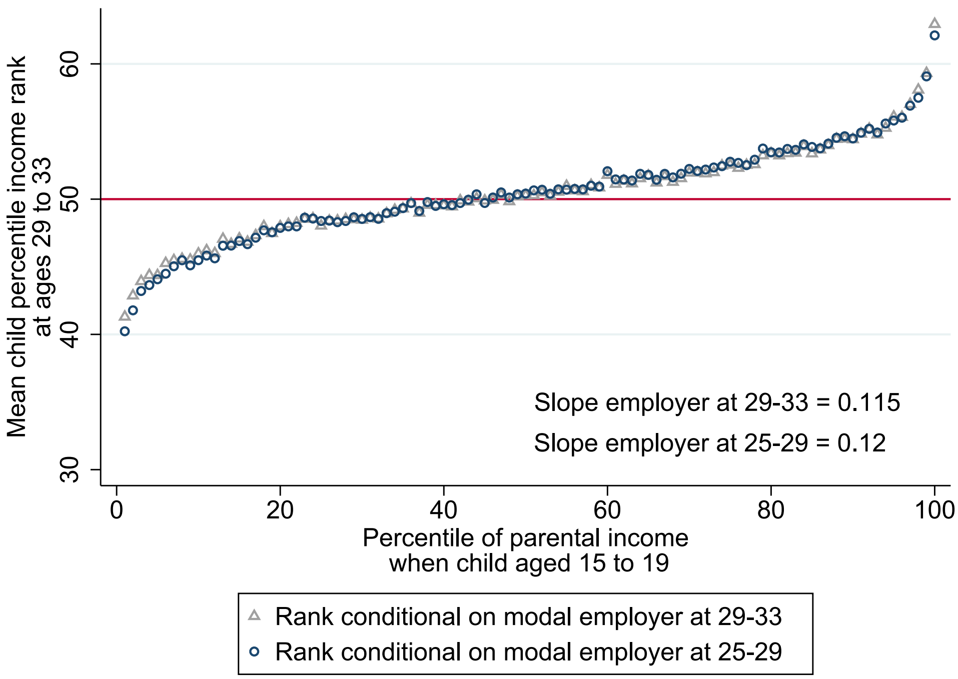

Notes: This figure shows mean child percentile income rank at age 29-33 for the 1984-86 birth cohorts. Both series are conditional on one's main employer. The series in grey triangle conditions on main employer at 29-33, and the series in blue circles conditions on main employer at 25-29.

Figure A8: Intergenerational Employment Income Mobility and Children's Employers

Panel A: Child-Parent Rank-rank Relationship

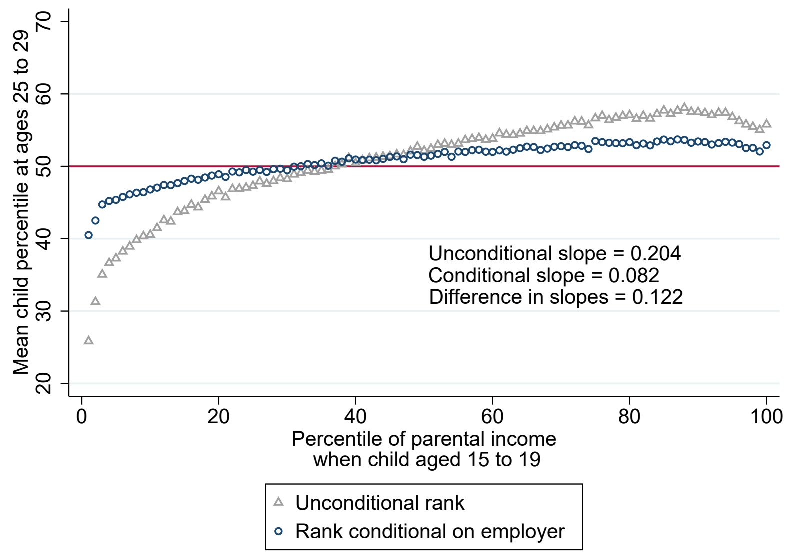

Panel B: Firm Sorting Component

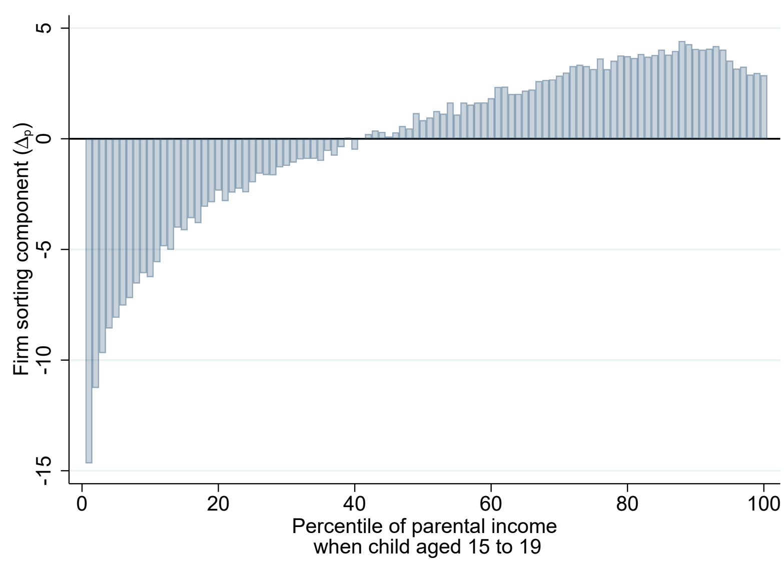

Notes: Panel A shows mean child employment income percentiles for each parental income decile. Grey triangles show unconditional means, whereas blue circles indicate conditional means accounting for differences in employers, as per equation (1). The difference in linear slopes represents the contribution of employers to employment income mobility. Panel B shows the firm sorting component. This component is equal to the vertical distance between the two series presented in panel A.

Figure A9: Distributions of Estimates of Social Connections Effects on Hiring Probabilities, With and Without Post-secondary Education

Panel A: Parent works at the firm

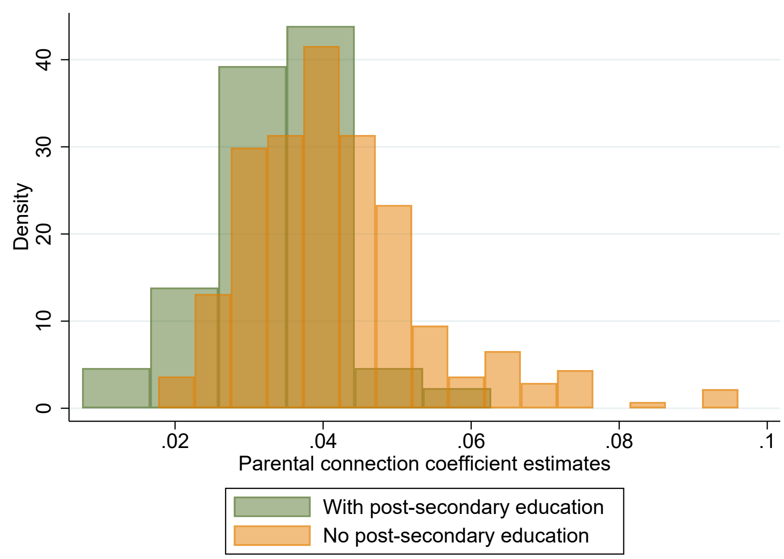

Panel B: Parent owns the firm

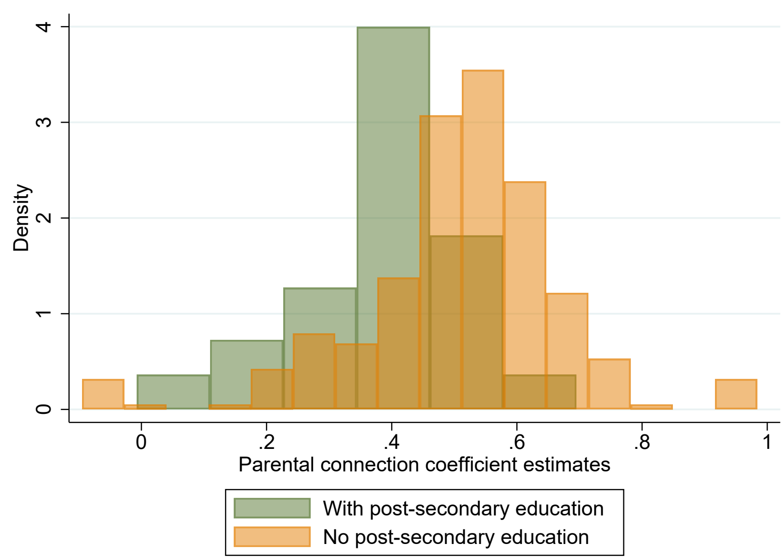

Notes:

A Data Appendix

A.1 Child-Parent Linkages

To create child-parent linkages, we mimic the methodology used by Statistics Canada in the creation of the Intergenerational Income Database (Corak and Heisz, 1999).

The CEEDD includes T1 Family Files (T1FFs) from 2001 onward. The T1FFs include family-level information for tax filers, their spouse, and their census family. We find all 15 to 19 year-olds (children) in T1FFs, and link them to their parents using unique census family identifiers. That is, for each child and each tax year, we identify who are the parents in their census family in that year. We then create time-invariant child-parents links using the earliest year a match is made. This means that if a child has different parents at ages 15 and 19, we create a child-parent linkage based on the census family they were part of when they were 15. In the process, we also impose additional restrictions, such as parents must be at least 16 years older than their children. This is mainly to avoid mistaking a child's relatively older roommate for a parent. Also, we drop any year in which a child has a spouse.

Some limitations must be noted. First, the parents we assign children may not be the ones they grew up with in early childhood. Rather, the parents are household heads in teenage years (15 to 19 years of age). Second, links can only be made for children who filed taxes at some point while still living with their parents. Despite these restrictions, the original IID, which was constructed using the same method we apply, has excellent coverage of the underlying population (Connolly et al., 2019). Reassuringly, we find that the income rank-rank relationship in our sample is very similar the one obtained using the IID (see Figure A4).

A.2 Employers

Employer identifiers are based on tax filing units.

When we assign each child their modal employer, if a child has two modes, we pick the employer from which they earned the most.

A.3 Education Groups

We create groups as unique combinations of a 4-digit CIP code, a credential type, a program type, a post-secondary institution, and an indicator for having graduated from the program. When institutions have more than one campus, we treat different campuses as different institutions. We merge small groups – those with less than 10 students in our sample of children born in 1987-89 – into broader categories, pooling together CIP codes (within a 2-digit category) within cells defined by a credential type, a program type, a post-secondary institution, and an indicator for having graduated from the program.

For children who have never attended a post-secondary institution in Canada (roughly a third of our sample of children), we define groups as census subdivisions based on their first recorded place of residence in the tax files. Census subdivisions correspond to municipalities, or to areas treated as municipal equivalents for statistical purposes by Statistics Canada. As such, they likely capture well the set of high schools these children attended. There are over 5,000 census subdivisions in Canada.

B Estimation Notes

B.1 Retrieving education-firm match effects

For each broad education category m, the estimating equation for the *full* dyadic sample is

$$H_{il} = \alpha_{e(i),l} + (\kappa^m + \gamma^m C_{il}) + \epsilon_{il} \tag{8}$$

where we normalize estimates of $\alpha_{e(i),l}$ such that their within-sample average is zero, and include an intercept κ_m. Since estimation of eq. (8) on the full sample is computationally infeasible, we must back-out the key parameters using estimates of γ^m as well as other moments of the data.

To back out the value of the intercept κ_m, note that by taking the group-m sample average

$$\kappa^m = \bar{H}^m - \gamma^m \bar{C}^m = \left(\frac{1}{J} - \gamma^m \bar{C}^m\right) \tag{9}$$

where J is the number of firms in the data, and $\bar{C}^m$ is the group-m sample average of C_{il}.

Then, averaging equation (8) by education-firm cell $\{e, l\}$ we obtain

$$\bar{H}_{e,l} = \alpha_{e,l} + \kappa^{m(e)} + \gamma^{m(e)} \bar{C}_{e,l}. \tag{10}$$

Note that $\bar{H}_{e,l}$ is the same as $s_{l|e}$ – the fraction of workers with education e that actually works at firm l – which can be directly measured. Combining equations (9) and (10),

$$\alpha_{e,l} = \left(s_{l|e} - \frac{1}{J}\right) - \gamma^m \left(\bar{C}_{e,l} - \bar{C}^m\right) \tag{11}$$

In education-firm cells with $\bar{C}_{e,l} = 0$ and $s_{l|e} = 0$, $\alpha_{e,l} = -\kappa_m$. With estimates of $\alpha_{e,l}$, κ^m, $\bar{C}^m$ and γ^m in hand, we can then compute predicted values $\tilde{H}_{il}$ for the entire dyadic sample

$$\begin{aligned}\tilde{H}_{il} &= \kappa^{m(e(i))} + \alpha_{e(i),l} + \gamma^{m(e(i))} C_{il} \\ &= \left(\alpha_{e(i),l} + \frac{1}{J}\right) + \gamma^{m(e(i))} \left(C_{il} - \bar{C}^{m(e(i))}\right)\end{aligned} \tag{12}$$

and average predictions by parental income

$$\tilde{s}_{l|p} = \sum_e \left[\underbrace{\left(\alpha_{e,l} + \frac{1}{J}\right) + \gamma^{m(e)} \left(\bar{C}_{ep,l} - \bar{C}^{m(e)}\right)}_{\tilde{s}_{l|e,p}} \right] \nu_{e|p} \tag{13}$$

where $\bar{C}_{ep,l}$ is the average of C_{il} for individuals in education group e and parental income group p, and $\nu_{e|p}$ is the share of children with parental income p who have education level e. In education-firm cells with $\bar{C}_{e,l} = 0$ and $s_{l|e} = 0$ (i.e. not a single member of group e works at or is connected to firm l), $\tilde{s}_{l|e,p} = 0$.

B.2 Sorting by comparative advantage

The within-firm component of the rank-rank income relationship can be decomposed

$$\beta_p = \sum_j \left(\bar{y}_{j|p} - \delta_j\right) s_j + \sum_j \left(\bar{y}_{j|p} - \delta_j\right) \left(s_{j|p} - s_j\right) \tag{14}$$

It is worth noting that when we break down β_p this way, the shares s_j may not add up to one *within the summations* anymore. That is, while $\sum_j s_{j|p}$ always sums up to one within each group p, the elements of s_j that are part of the sum $\sum_j \left(\bar{y}_{j|p} - \delta_j\right) s_j$ won't sum up to one. This is because workers of a given income group p are not observed at all firms, and so $\bar{y}_{j|p}$ is non-existent in many cases.

Since we do not know what the average income of group p would be at firms at which they are not observed, it is difficult to assess the extent to which groups of workers sort into firms at which they experience an income advantage relative to other groups. We therefore rely on sorting patterns within the set of observed matches to evaluate whether group-based sorting on the basis of comparative advantage is important. The idea is that sorting patterns at the "intensive margin" (i.e. more vs fewer workers of a given group are hired by firm j) are informative of sorting patterns at the extensive margin (i.e. some vs no worker of a given group is hired by firm j).

First, we calculate the correlation between $(\bar{y}_{j|p} - \delta_j)$ and $(s_{j|p} - s_j)$, and find that in our main specification it is -0.04. If anything, workers from groups that earn the most among all workers at firm j are under-represented at that firm. Second, we compare the true

$\beta_p = \sum_j \left(\bar{y}_{j|p} - \delta_j\right) s_{j|p}$ with a properly re-scaled sum that uses the population-wide shares s_j as weights. That is, we calculate for each group p:

$$\tilde{\beta}_p = \frac{\sum_j \left(\bar{y}_{j|p} - \delta_j\right) s_j}{\sum_j 1\{s_{j|p} > 0\} s_j}. \tag{15}$$

To ease the interpretation of $\tilde{\beta}_p$, let $\beta_p^{complete}$ denote the (unobserved) value of β_p one would obtain if group p was observed at all firms in the data. $\tilde{\beta}_p$ will coincide with $\beta_p^{complete}$ if the counterfactual average within-firm gap for group p at firms where they are not observed was the same as the average within-firm gap at firms where they are observed.

In Figure A10, we plot the re-scaled sum from eq. (15) against the true β_p. The two series closely overlap throughout the parental income distribution, expect for the top percentile, where the re-scaled sum slightly mutes the non-linearity.

Figure A10:

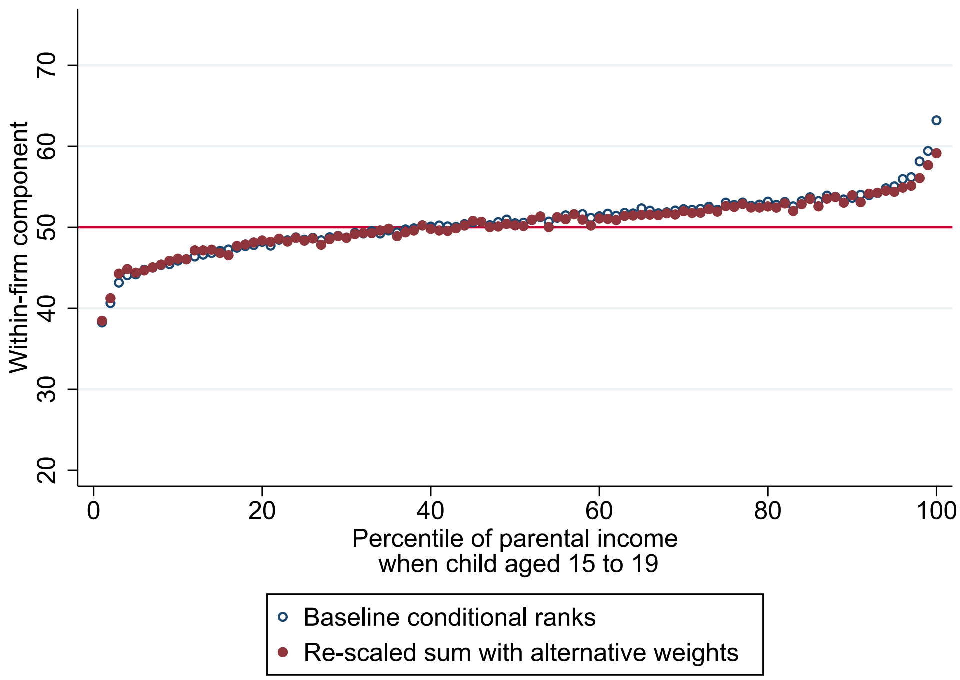

Notes: This figure shows

B.3 Estimation of AKM worker and firm effects

We start with a sample that includes all 24-54 year-old workers in Canada for years 2001 to 2018, and impose a few sample restrictions. First, we only keep workers with annual earnings above 5,000$. We then drop workers who appear in the data for less than three years. Finally, we focus on the largest connected set (Card et al., 2018).

The dependent variable is the log of T4 earnings for worker i in year t. We remain agnostic regarding the source of firm premiums, whether it is a wage or a work hours effect. A usual AKM specification takes the form

$$\ln e_{it} = \alpha_i' + \psi_{j(i,t)}' + \beta X_{it} + \varepsilon_{it}' \tag{16}$$

where X_{it} is a vector of covariates, which generally include year fixed effects and a polynomial in age. Unfortunately, the server on which we estimate the model is unable to estimate such a specification on the full sample, running into memory errors. To circumvent this computational issue, we first regress $\ln e_{it}$ on a full set of year and age fixed effects, and save the residuals r_{it}. We then regress these residuals on worker and firm fixed effects

$$r_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it} \tag{17}$$

For validation purposes, we estimated eq. (16) on a subsample of workers and firms in Ontario (Canada's largest province), and compared one-step estimates to those obtained using the two-step approach described above. The correlation in person effects is 0.965, and the correlation in firm effects is 0.999.

B.4 Selection-adjusted income-rank firm premiums

Let I_i denote child i's average income between the age of 25 and 29, and $\psi_{j(i)}$ is the estimated AKM firm effects for child i's modal employer j between the age of 25 and 29. We can then write

$$\ln I_i = a_i + \psi_{j(i)} \tag{18}$$

where a_i implicitly captures all components of child i's income that aren't due to firm effects. With data on both $\ln I_i$ and $\psi_{j(i)}$, we can back out values of a_i. Then, we calculate a counterfactual income as

$$\tilde{I}_i = \exp(a_i + \bar{\psi}^c) \tag{19}$$

where $\bar{\psi}^c$ is the sample average of $\psi_{j(i)}$, which we calculate separately by birth cohort c. Then, using cohort-specific income percentile cutoff points from the true income distributions, we retrieve the income rank $\tilde{y}_i$ for counterfactual income $\tilde{I}_i$. For each child, the difference between their real and counterfactual income rank $y_i - \tilde{y}_i$ represents the contribution of the firm premium to their income rank. We finally calculate average firm premiums for firm j as $E[y_i - \tilde{y}_i|j]$. This corresponds to the average contribution of firm j to the income rank of its workers.