NETWORK
LISER

Discussion Paper Series

IZA DP No. 18691 May 2026

Social Connections and the Persistence of Income Across Generations

Jean-William Laliberté
University of Calgary
and IZA@LISER

Alexander Whalley
University of Calgary

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18691 May 2026

Social Connections and the Persistence of Income Across Generations*

Abstract

We use matched parent-child-employer-employee data from Canada, linked to detailed educational records, to quantify the contribution of social connections to employers to intergenerational income mobility. Sorting across employers accounts for roughly a third of the transmission of income across generations. To estimate the impact of social connections on differential representation across employers, we compare classmates – those with the same degree from the same institution – who have different social connections. We find social connections in the labor market explain about 15% of the firm-sorting component of the intergenerational income rank-rank relationship, about a third the explanatory power of education.

JEL classification

J62, J31, J24, L25, E24

Keywords

social connections, intergenerational mobility

Corresponding author

Jean-William Laliberté

jeanwilliam.lalibert@ucalgary.ca

* We are grateful to Nathanael Hammond for stellar research assistance. We also thank Pamela Campa, Raj Chetty, Laura Hering, Henrik Kleven, David Margolis, Olivier Marie, Sacha Kapoor, Erik Plug, Abhijeet Singh, Matt Staiger, Jan Stuhler, Katarina Wessling, Dean Yang, Tom Zohar, and numerous seminar and conference participants for helpful conversations. We also thank the Canadian Research Data Centre Network for support.

"*The income of children is raised when they receive more human and nonhuman capital from their parents. Their income is also raised by their endowment of genetically determined race, ability, and other characteristics, family reputation and connection and knowledge, skills and goals provided by the family environment.*" Becker and Tomes (1979)

1 Introduction

The persistence of income across generations is an empirical regularity across many economies. The human capital forces emphasized by Becker and Tomes (1979, 1986) have long been central to understanding this phenomenon. Yet, the quote in the epigraph also emphasizes the importance of non-human capital forces such as family reputation and connections. Despite the prominence of social contacts in hiring (Topa, 2019; Granovetter, 1995) and the persistence of family dynasties – Ford, Mars, and Murdoch families for example – their quantitative importance for intergenerational income persistence remains unclear. We study how social connections to employers affect intergenerational income mobility to fill this gap.

Estimating the effect of social connections on children's job opportunities is challenging. First, worker sorting across employers is not random. Preferences and skills are likely correlated across social contacts. That is, two individuals connected to each other, such as a parent and their child, likely share attributes that make them both likely to work for the same firm independently of the fact that they know each other. They may search in the same local labor market or may pursue similar careers. Second, few datasets include information on the identity of workers, their employers, as well as their social contacts, which is necessary to estimate social connection effects.

We use newly linked data from Canada, matching administrative tax data that include both employee-employer and child-parent linkages, with highly granular post-secondary education enrollment and graduation records, to address these issues. Isolating the effect of social connections on job opportunities requires building counterfactual worker-job match probabilities in the absence of social connections. To do so, we apply a "classmates" research design that narrowly compares the probability of working at a given firm between people who have the exact same degree but not the same connections. The highly detailed nature of our education data – where we know if a child completed a Master's degree in Computer Science at the University of Waterloo, or a certificate in Pastry Arts at Vancouver Community College, for example – helps us isolate the influence of connections on job match probabilities from other confounding mechanisms such as parents and children making similar educational investments or living in similar areas.

Our exploration of the role of social connections to employers for intergenerational income mobility proceeds in several steps. The starting point is the intergenerational rank-rank income relationship (Chetty et al., 2014), which effectively partitions the sample of children into 100 mutually exclusive parental income groups. In a first step, we quantify the role sorting to firms play, in an accounting sense, for the persistence of income across generations.[1] This opening analysis characterizes the differences in firm sorting behavior across parental income groups we seek to understand. We find that differences in representation at high-paying firms across parental income groups – the between-firm sorting component – accounts for close to a third of the intergenerational income rank-rank relationship.

Our main analysis quantifies how social connections influence the sorting of young workers across firms. Our approach isolates the effect of social connections on job opportunities by building counterfactual worker-job match probabilities in the absence of social connections. That is, we identify the impact of connections on match probabilities by comparing people who enrolled in *the same* educational program and therefore faced similar labor market prospects, as in Kramarz and Skans (2014), but have *different* social connections.[2] Concretely, we estimate the impact of social connections on job match probabilities between each worker and each firm in the data, and use these estimates to produce counterfactual allocations of workers across firms.

Although a person's entire social network is ultimately unobservable to the analyst, we cast a broad net and consider numerous connection types, including both family ties and non-kin relationships. For instance, it is well-established that many adult children work at firms that also hire one of their parents (Staiger, 2025; Corak and Piraino, 2011). Our analyses therefore include such connections to a parent's employer. For all private incorporated firms, the data also include the list of business owners. With that information in hand, we can separately identify family firms. For example, among children in the top 1 percentile of parental income, more than one child out of ten work *for* their parents in their late 20s.[3]

[1]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.

[2]Kramarz and Skans (2014) use Swedish graduation records. Our database includes both graduation and enrollment, and therefore allows us to further differentiate between people who started but never graduated from a post-secondary program, and those who did complete it. Given relatively low completion rates in Canada (Childs et al., 2017), as is the case in the US, and known differences in economic outcomes between workers with some college education and those with earned degrees (Oreopoulos and Petronijevic, 2013), distinguishing between graduates and non-completers is important.

[3]The share is considerably higher at younger ages, but we focus on employment between the ages of 25 to 29 in our main analyses.

We additionally consider social connections to firms via siblings, parents' former co-workers, and classmates from one's graduation cohort.[4]

The impact of social connections on job match probabilities differs considerably in magnitude across connection types. For instance, while the likelihood of working at a given firm increases by about 3 to 5 percentage points when one's parent works there, the effect size is a full order of magnitude larger in cases where a parent is the owner of the firm. In that latter case, most estimates range between 28 and 65 percentage points increases. The probability of matching with an employer increases by around 2 percentage points if a sibling works there, and by close to 7 percentage points if a sibling owns the firm. Finally, the impact of non-kin ties is considerably smaller: the presence of a parent's past co-worker on a firm's payroll increases the probability of matching with that firm by close to 0.4 percentage point. If a parent's past co-worker owns a business, the likelihood of working at this firm rises by 0.1 percentage point. Whether a same-cohort classmate's parent works at or owns a firm has virtually no impact on the probability of working there. Despite our granular controls for education, if children and their social contacts share unobserved firm-specific skills or preferences, our estimates may reflect matching based on correlated unobserved skills or preferences rather than social connections. To verify whether this is the case, we show that social connections in the future – based on matches that have not yet occurred – do not affect sorting today. This rules out time-invariant firm-specific preferences or skills driving our results.

How economically significant are social connection effects for intergenerational income mobility? The quantitative importance of connections depends not just on effect sizes but also on which firms children from different backgrounds are connected to. Combining both pieces of information, we find that connections to a parent's employer are the most important ones for income mobility. They explain close to 10% of the between firm sorting component. In contrast, connections to parents' family firms (i.e., businesses owned by parents) explain less than 1% of the firm sorting component. This is partly due to the fact that few parents own businesses, and those who do are dramatically over-represented at the top of the income distribution. Hence, family firms have little explanatory power for the transmission of income across generations, except at the very top. Connections to siblings' employers and to parents' past co-workers' employers respectively account for 3% and 1% of the firm sorting component. The explanatory power of connections through same-cohort classmates' parents is essentially

[4] We focus on connections through same-cohort classmates because our main econometric specification includes education-by-firm fixed effects, leaving no residual variation in social connections shared by all students (of any cohort) of a given education program.

null. Overall, social connections collectively explain for almost 15% of differences in the firm sorting component across parental income groups, although this figure is likely a lower bound since we cannot observe all possible connection types.

To put these numbers in perspective, consider that the share of the between-firm sorting component that can be explained by differences in education is 45%. Hence, the contribution of social connections is about a third that of observable differences in education. The share attributable to education captures a causal effect of education on job matching probabilities, but also the influence of other correlated attributes such as a school's proximity to a firm, differences in academic ability, and, importantly, social connections built in school. That is, a fraction of the component we attribute to education may actually reflect social connection effects via faculty members, alumni networks, and participants in co-op programs, for example. This reinforces our interpretation of the share of the sorting component we attribute to social connections as a lower bound.[5]

Why do firms hire socially connected workers? We consider several classes of explanations. Firms may seek to hire connected workers because social connections provide information about the quality of the match or alignment with the firm. This matching channel would lead connected workers to be paid more than their coworkers and increase labor productivity at the firm. Alternatively, hiring of connected workers may reflect preferences of workers for firms – e.g. if workers enjoy working with people they know – and preferences of firms for workers – e.g. favoritism. To guide the interpretation of our findings and evaluate which of these mechanisms is most likely at play, we conduct auxiliary analyses in which we (i) compare the earnings of socially connected workers to that of their unconnected co-workers, and (ii) use firm balance sheet data to examine how firm performance changes when a socially connected worker is hired.

We find that workers who have a social connection to their employer earn more, on average, than their co-workers who do not have such a connection. This is true of all connection types, but the income gaps between connected and unconnected co-workers is the largest in cases where the connected worker is either the child or the sibling of the firm owner. Perhaps surprisingly, the income gaps are roughly the same for connections to a social contact's employer, regardless of the contact type. Are connected workers relatively more productive, or do they benefit from preferential treatment? To distinguish between these

[5]In Appendix D, we provide indirect evidence of education-based connection effects by netting out the contribution of aggregate factors that are associated with geography and fields of study, and are shared too widely in the population to plausibly capture relevant social networks. We find the remaining residual education program-firm match effects explain 5% of the firm-sorting component.

interpretations, we examine changes in firm performance around the time a connected worker joins the firm. Here, we find that firm revenue per worker and value-added per worker decline substantially after a connected child joins a family firm (i.e., a firm their parents own), and there is no evidence that the firm eventually makes up for this loss over a 5-year horizon. A decline in revenue per worker is also observed when a child joins their parent's employer, although the magnitude of the change is much smaller than for family firms. We find no change in firm performance on the basis of other connection types. Overall, these results suggests within-firm income gaps between connected and unconnected workers are unlikely to reflect differences in labor productivity.

Putting all of these findings together, we conclude that differences in human capital cannot be the sole factor driving the persistence of income across generations. Social connections do allow children from high-income families to have disproportionate access to high-paying firms. In some cases, this might be because their social connections allow them to signal their relatively high productivity, something comparable high-productivity workers from lower-income families may not have the opportunity to do. Yet, we also find suggestive evidence that children who have access to jobs at family firms may benefit from favoritism.

We see our study as making four contributions. First, we contribute to a vast literature on intergenerational income mobility (Becker and Tomes, 1979; Solon, 1999; Sacerdote, 2007; Chetty et al., 2014; Stuhler, 2018). Most investigations of intergenerational income transmission mechanisms focus on differences in human capital that can be traced back to differences in either inherited attributes (Björklund et al., 2006) or parental investments in childhood (Agostinelli and Wiswall, 2025; Caucutt et al., 2020). Other examples of pre-market factors include parenting practices (Doepke et al., 2019), neighborhoods (Chyn and Katz, 2021), and colleges (Chetty et al., 2020). There is a smaller but growing 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. (2025)). Our work shows that social connections to early career employers are important for intergenerational mobility, with their explanatory power being roughly a third that of observable differences in detailed education.

Second, we contribute new evidence on how social factors affect inequality. Theoretically, hiring through social connections can cause one group's (Calvo-Armengol and Jackson, 2004) or one generation's (Bolte et al., 2020) employment advantage to persist. Empirically, referred workers perform better and stay longer at hiring firms (Dustmann et al., 2016; Burks et al., 2015; Pallais and Sands, 2016), though whether referrals reduce or amplify income inequality remains unclear given heterogeneous effects across worker skill levels (Brown et al.,

2016) and connection types (Lester et al., 2023). Eliason et al. (2023) find that network-based matching reduces inequality, while Chetty et al. (2022) show that social capital correlates positively with intergenerational mobility across places. We provide new evidence that social connections can instead *hinder* intergenerational income mobility – specifically, connections *to firms* operating through the labor market play a very different role from social capital *in a location*, which can operate through non-market channels. This complements Staiger (2025), who finds large positive effects of working at a parent's employer, and San (2022), who shows social connections amplify ethnic wage differentials in Israel. Our work extends this literature in three ways: we construct complete counterfactual worker-firm allocations to capture connections to *potential* employers beyond realized matches, consider a broader set of connection types including novel evidence on connections to business owners, and provide new evidence on how connections affect firm performance.

Third, we contribute to understanding the role firms play in inequality. Seminal contributions quantify how firms drive inequality trends (Card et al., 2013; Song et al., 2019), as well as inequality by gender (Card et al., 2016), race (Gerard et al., 2021), and immigration status (Dostie et al., 2023). Recent papers show that sorting across firms meaningfully shapes the intergenerational income elasticity (IGE) in Israel and Sweden (Dobbin and Zohar, 2025; Engzell and Wilmers, 2025; Forsberg et al., 2026). Our results on firms' contribution to mobility are quantitatively similar to theirs, suggesting these findings generalize across countries and measurement approaches.[6] Our distinct contribution is to open the black box of *why* children from higher-income families sort to better firms. Building on the classmates research design of Kramarz and Skans (2014), we show that social connections account for a non-negligible share of the income gradient in firm sorting – documenting a mechanism largely absent from prior work.

Last, we contribute to the literature on family firms. While a large fraction of businesses throughout the world are organized around families (Burkart et al., 2003), why this is so remains subject to debate (Bertrand and Schoar, 2006). One strand of research has shown that family based CEO succession reduces firm performance (Bennedsen et al., 2007). Yet, evidence on wage differentials indicates that retaining private benefits of control appears to be valuable (Di Porto et al., 2024). We provide evidence that hiring decisions in family firms are likely partly based on non-financial considerations, since the businesses who hire

[6] Each of these papers decomposes the IGE into individual and firm components using an AKM-type two-way fixed effects specification (Abowd et al., 1999). Dobbin and Zohar (2025) find that AKM firm effects account for 22% of the IGE in Israel; Engzell and Wilmers (2025) and Forsberg et al. (2026) find shares of 28% and 36% in Sweden. We find selection-adjusted firm effects account for 31% of the rank-rank slope in Canada.

the owner's children do not seem to benefit economically from doing so. Still, we find the impact of family firms on intergenerational mobility is modest, except at the very top of the distribution.

The remainder of the paper is organized as follows. The data are described in Section 2. Section 3 presents a decomposition of the rank-rank relationship into within- and between-firm components. Section 4 estimates the impact of social connections on job match probabilities and quantifies their importance for income mobility. Section 5 probes alternative explanations for why social connections affect job match probabilities, and Section 6 concludes.

2 Data

We use administrative tax data from the Canadian Employer Employee Dynamics Database (CEEDD), linked with granular education records from 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, hospitals 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), for each fiscal year we match workers with the employer from which they earned the most within that year.

The PSIS records individual-level enrollment and graduation at all public and private not-for-profit post-secondary institutions in Canada. This covers virtually all post-secondary students in Canada. The database notably includes information on which institution a student is enrolled in, the program and credential types, as well as the field of study, coded

[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.

using 6-digit Classification of Instructional Programs (CIP) codes.

2.1 Intergenerational Sample

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 (i.e., 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 know their education. In contrast, coverage of educational outcomes for later birth cohorts is excellent, but we don't observe these cohorts' income at older ages. Our choice of cohorts is meant to balance the trade-off between coverage of education outcomes and coverage of income later in life. 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.

2.2 Variable Definitions

Income ranks. Our definition of income variables largely follows prior work on intergenerational income mobility in Canada (Haeck and Laliberté, 2025; 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 income of children as the average between the ages of 25 and 29 inclusively to reduce the influence

[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 impute their income using their unreported T4 earnings.

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

Education groups. From PSIS, we define an education group e as a unique combinations of a field of study (defined as a 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. This classification is very granular as there are over 450 different fields of study.[10] For workers with no post-secondary education, we define an education group as 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.

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.[11] Similarly, prior work on parental connections generally focus on children's first stable job (Staiger, 2025; San, 2022; Kramarz and Skans, 2014). This procedure assigns children in our estimation sample to over 230,000 different employers.

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 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,

[10]For example, there are 41 different fields related to engineering, such as civil engineering, chemical engineering, materials engineering, nuclear engineering and petroleum engineering.

[11]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 or in entrepreneurship (Gendron-Carrier, 2023). Differences in income across firms should therefore be interpreted as reduced-form income gaps between people who have different main employers.

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). 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. We discuss the implications of the employment margin for intergenerational income mobility in the next section.

Social connections. We consider four categories of social contacts (parents, siblings, parents' past co-workers, and parents of graduation-cohort classmates), and two kinds of relationship to a firm (either working at or owning the firm), resulting in eight types q of social connections for which there is identifying variation within education groups. We define the corresponding social connection variables C^q_{il} as indicators for whether child i is socially connected to firm l. For instance, for connections to a parent's employer, C^q_{il} takes a value of one if at least one of child i's parents was employed by firm l at any point within the 5-year window we use to assign children's main employer. Similarly, the indicator variable for connections to firms parents of child i own takes a value of one if at least one of child i's parents owned firm l at some point within the 5-year window. Connections via siblings are similarly defined.

In terms of non-family ties, parents' past co-workers are defined as individuals who worked together with the parent (i.e., for the same employer at the same time) at some point when child i was between the age of 20 and 24, and no longer works with the parent (i.e., did not share an employer with the parent when the child was 25-29). We only keep co-worker relationships from firms with less than 100 employees to make sure the co-workers plausibly knew each other, and exclude contemporary connections to a parent's previous employer.

Finally, we include connections through post-secondary education classmates from one's graduation cohort. We rely on variation across cohorts because our main specification in Section 4 includes a full set of education-by-firm fixed effects, thereby absorbing any social connection shared by all graduates of a given program. That is, we can only identify the impact of differences in social connections within education programs.[12] Although students in different cohorts of a given program may interact with each other, we assume two students are more likely to know each other if they graduated at the same time.[13] We therefore partition

[12]Using a similar econometric specification, Fischer et al. (2025) leverage random assignment to tutorial groups within a business program in Copenhagen.

[13]Kirkebøen et al. (2021) find the vast majority of spousal matches occurring within education programs in Norway are within cohorts of applicants.

Figure 1: Social connectedness of employees, by parental income rank

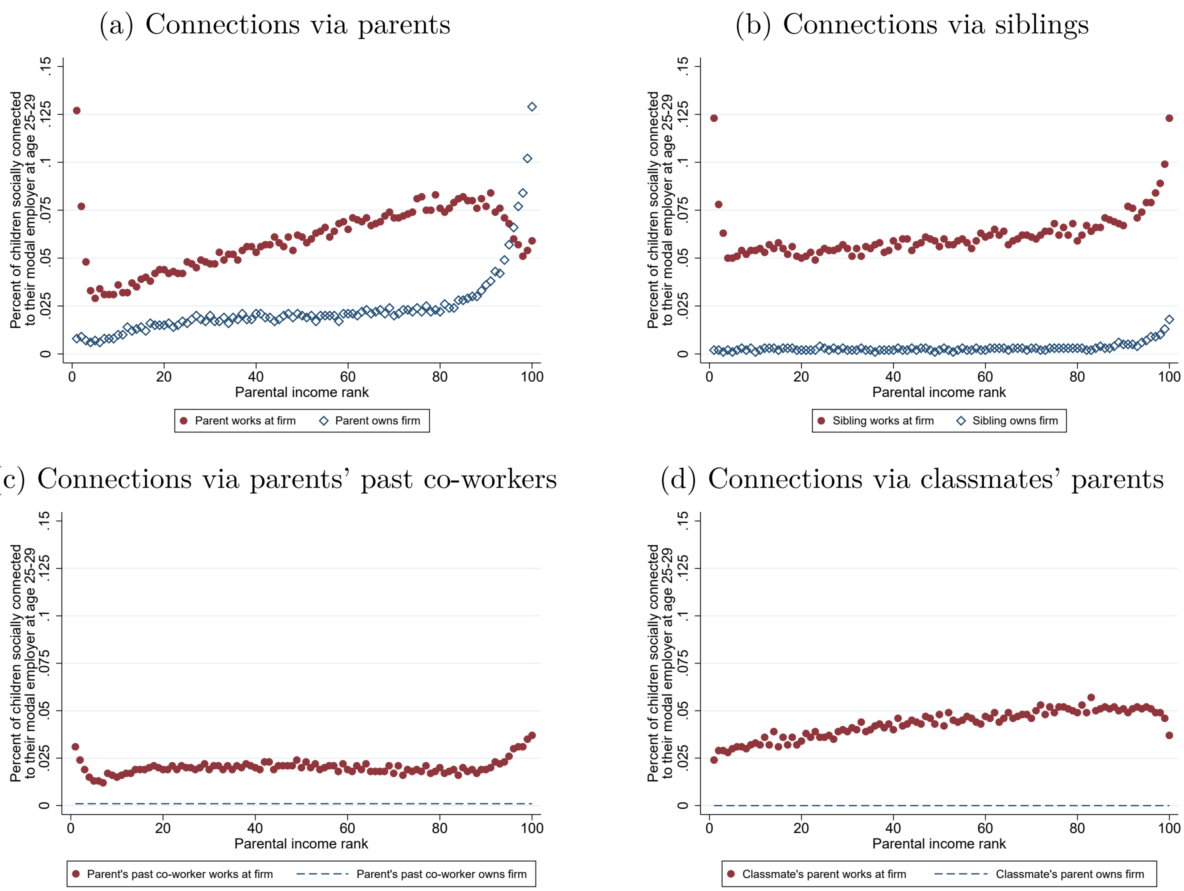

Notes: This figure shows the fraction of children in each percentile of the parental income distribution who have a social connection to their modal employer between the ages of 25-29. To construct these statistics, children with no modal employer are excluded. In each panel, the fraction of children who share their contact' employer excludes cases where the contact also owns the firm. In panels (c) and (d), the parental income-specific shares of children working at firms their social contacts own are too small to be released per vetting guidelines. Instead, the figure shows the overall shares (pooling all parental income groups) in those cases. In panel (d), connections to classmates' parents only include classmates from one's graduation cohort.

education groups e into graduation cohorts based on graduation years, and code indicators C^q_{il} for whether a parent of a child in that cohort-education cell (excluding one's own parents) works at or owns firm l. We focus on this indirect link to a classmate's parent rather than a direct link to the classmate to ensure our estimates are not driven by cross-cohort differences in human capital and market conditions. When generating the set of same-cohort peers, we include graduates from birth cohorts outside our focal sample of children.[14] Finally, we only keep connections for program-specific graduation cohorts with at most 100 graduates to ensure students knew each other.

How do social connections relate to parental income? In Figure 1 we report the fraction of working children who are socially connected to their modal employer, separately by parental income rank and connection type. The fraction of workers whose main employer is also their parent's employer is 6 percent, on average.[15] This share is increasing in parental income between the 5th and 90th percentiles, a result potentially partly driven by a positive relationship between parental income and the likelihood of having two working parents (and therefore connections to more firms).[16] The probability a child's modal employer is their parent's employer is much higher for the bottom 2 percentiles, and considerably lower at the top of the distribution.[17] In contrast, the likelihood of working at a firm one's parents own is almost flat (around 2 percent) between the 20th and 80th percentiles, and increases precipitously above the 90th percentile. 13 percent of working children in the top 1 percent of parental income work *for* their parents.[18]

Roughly 6 percent of working children share their sibling's employer, and only 0.3 percent work at a firm their sibling owns. Family ties between siblings seem particularly important at the top of the parental income distribution. About 2 percent of working children are employed at a firm that also employs a former co-worker of one of their parents, and only 0.3 percent work at a business owned by their parents' former co-worker. Finally, 4 percent of

[14] We include connections to graduates born in 1984-86 and 1990-92 for whom family ties are known. This means the set of peers excludes international students whose parents don't live in Canada.

[15] Most business owners pay themselves a salary, which means they are also considered employees of the firm they own. To avoid double counting, the share of children working at firms where their social contact also works shown on the figure excludes cases where the social contact owns the firm.

[16] Figure A2 plots the share of children in lone-parent families by parental income rank. The overall fraction of children with a single parent in our data is just below 20%, which is relatively lower than the corresponding share of 30.6% in Chetty et al. (2014).

[17] The high share of children and parents/siblings sharing an employer at the very bottom of the income distribution is mostly explained by the combination of (i) an over-representation of Indigenous children at the bottom of the parental income distribution and (ii) an over-representation of Indigenous families with very low income among those working in local Indigenous Public Administration (Connolly and Haeck, 2026).

[18] This fraction only includes those whose parents' business is their modal employer between the age of 25 and 29. Many more have worked for their parents at some earlier point in their lives.

working children share an employer with one of their same-cohort classmate's parent. This share generally increases with parental income, in part because children from high-income families are more likely to graduate from post-secondary programs and so to have such connections. The share of children working at a firm owned by the parent of one of their same-cohort classmate is indistinguishable from zero.

Overall, these comparisons suggest a role for connections to employers in intergenerational mobility, but may partly capture correlated attributes rather than social connection effects. For instance, the fact that several children work at firms that also employ at least one parent of their same-cohort classmates likely reflects overlap in skills and geography, not a social connection effect.[19] In Section 4, we compare observed to potential matches to reach a more definitive conclusion. Figure 1 also highlights the value of our rank-rank approach as the relationships between parental income and social connections are non-linear with much variation in the tails of the distribution.

3 Intergenerational Income Mobility and Employers

3.1 Decomposition Method

To first evaluate the role of firm sorting for income mobility, we decompose the average income rank of children with parental income p, $\bar{y}_p$, into a component capturing sorting across employers by parental income rank (denoted by Δ_p), and one reflecting within-firm income persistence:

$$\bar{y}_p = \underbrace{\sum_j \delta_j s_{j|p}}_{\equiv \Delta_p} + \sum_j \left(\bar{y}_{jp} - \delta_j\right) s_{j|p} \tag{1}$$

where j indexes employers, δ_j is a measure of average compensation at employer j, $s_{j|p}$ is the share of children of group p who are employed at firm j, and $\bar{y}_{jp}$ is the average income rank of children in firm j belonging to income group p. This equation highlights that the firm sorting component Δ_p depends on two key sets of parameters. First, the distributions of employers within parental income groups ($s_{j|p}$) characterize the allocation of children to firms. Differences in these distributions across parental income groups indicate whether

[19] By construction, classmates share skills and location. Parents and children often do too.

children from different backgrounds are matched with different employers.[20] In the next section, we investigate how social connections shape these distributions. Second, dispersion in compensation δ_j across firms affects 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 δ_j is close to zero), then firm sorting by parental income would be inconsequential for income gaps.

The formulation of equation (1) is very general; we can plug-in any compensation structure to examine the sensitivity of our findings to measurement of firm-level compensation. For completeness, we implement the decomposition using two different measures of δ_j. The first one aims at documenting the total role of between-firm income inequality for mobility. For the purpose of this descriptive accounting exercise, we regress the income rank of child i on parental income rank dummies and a full set of firm fixed effect, and summarily evaluate the role of between-firm inequality by examining how much accounting for firms fixed effects flattens the relationship between children's and parents' income ranks.[21] Under this approach, the estimated δ_j's reflect firm-level average income ranks net of differential representation by parental income. The second measure of δ_j we consider isolates the role of firm pay premiums. That is, we "selection-adjust" firm effects for compositional differences in time-invariant unobserved earnings potential using typical AKM (Abowd et al., 1999) worker and firm effects obtained from the entire sample of prime-aged Canadian workers. For conciseness, we summarize the main results of this analysis below, but refer the reader to Appendix B.3 for the details of the procedure used to adjust firm effects for selection. Since unadjusted firm effects 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, a comparison of results across measures of δ_j informs patterns of selection into firms by earnings potential.

3.2 Results

Figure 2 plots the unconditional average income rank $\bar{y}_p$ alongside estimates of average ranks conditional on employers (i.e., the within-firm component). The linear slope of the unconditional relationship is 0.232, very similar to prior work using Canadian data (Connolly

[20] Differences in theses distributions may contribute to observed within-firm gaps if workers sort into firms on the basis of group-based comparative advantage. Appendix C.2 provides evidence that sorting by comparative advantage is limited in our setting.

[21] The firm fixed effects estimation procedure is described in detail in Appendix B.1.

et al., 2019; Connolly and Haeck, 2024; Haeck and Laliberté, 2025).[22] 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, 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 conditional series – which correspond to the within-firm component – exhibit similar non-linearities at the top end of the parental income distribution, indicating that these patterns are not explained by firm sorting. The income advantage of children growing up in top 1 percent families relative to top 10 percent families is therefore not due to differential representation at "superstar" firms. The shape of the conditional rank-rank relationships at the bottom of the parental income distribution is a bit more linear than the unconditional one, but a salient steepening of the slope remains, particularly below the 5th percentile.

Overall, the linear rank-rank slope conditional on firm fixed effects (not adjusting for selection) is 0.115, indicating that between-firm income inequality explains roughly half $((0.232 - 0.115)/0.232 = 50.4\%)$ of the transmission of income across generations. The corresponding slope estimate based selection-adjusted firm effects is 0.160, suggesting employer-specific pay premiums explain close to a third $((0.232 - 0.160)/0.232 = 31\%)$ of the transmission of income across generations. This number is strikingly close to comparable estimates: Dobbin and Zohar (2025) find that AKM firm effects account for 22% of the intergenerational income elasticity in Israel, whereas Forsberg et al. (2026) and Engzell and Wilmers (2025) find corresponding values of 36% and 28% in Sweden.[23]

Graphically, the firm sorting component Δ_p is equal to the vertical distance between the conditional and unconditional series in Figure 2, panel (a). For greater visual clarity, we plot the firm sorting component in panel (b). The benchmark series not adjusted for selection shows a sharp steepening below the 20th percentile of the parental income distribution, a pattern that is partly muted in the selection-adjusted series. That is, segregation of low-pay workers in low-pay firms partly drives the persistence of income at the bottom of the distribution. Interestingly, there appears to be slight patterns of *negative* selection into

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

[23] Engzell and Wilmers (2025) report a host of estimates for different subsamples and specifications. The 28% figures is obtained by adding the worker-to-firm and firm-to-firm components reported in their Table A6, focusing on sons. Forsberg et al. (2026) further demonstrate the share of the IGE attributable to firm effects rises to 50% once dynamic returns are taken into account.

Figure 2: Intergenerational Income Mobility and Children's Employers

(a) Child-Parent Income Rank-rank Relationship

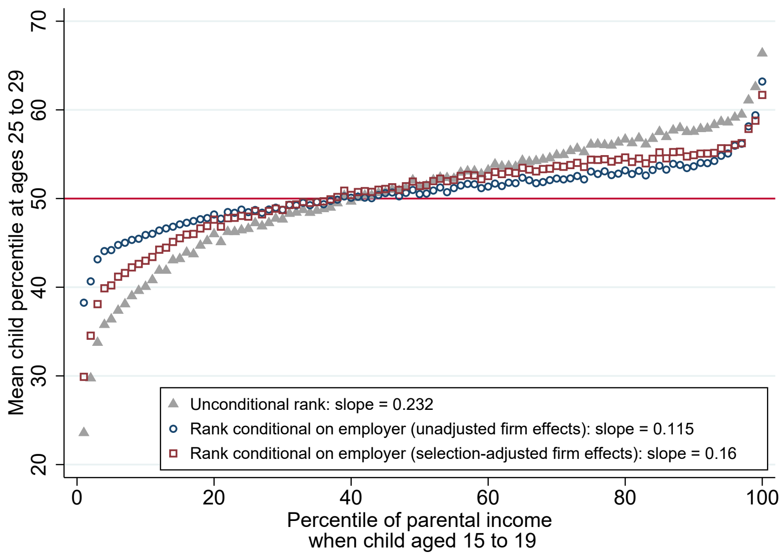

(b) Firm Sorting Component

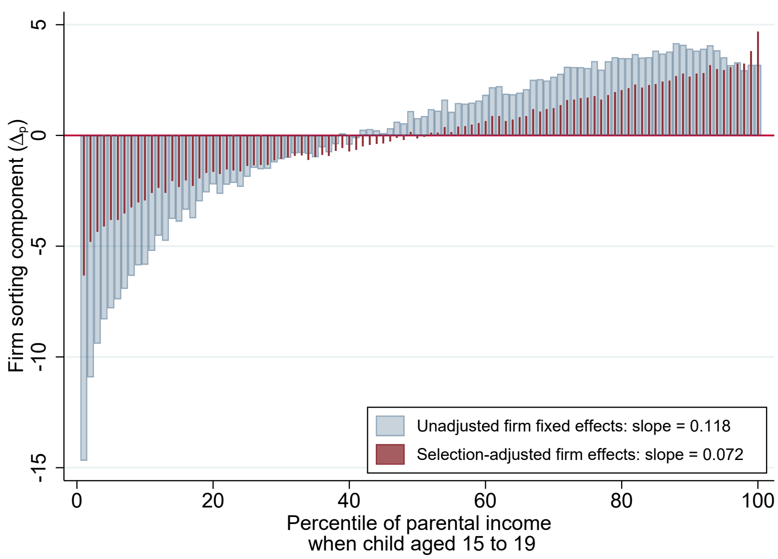

Notes: Panel (a) shows mean child income percentiles for each parental income decile. Grey triangles show unconditional means, whereas blue circles and red squares indicate conditional means accounting for differences in employers using unadjusted and selection-adjusted firm effects, respectively. Panel (b) shows the firm sorting component. This component is equal to the vertical distance between the conditional and unconditional series presented in panel (a).

firms at the very top of the parental income distribution, as the selection-adjusted sorting components are larger in magnitude than the unadjusted ones. In other words, low-pay workers from rich families disproportionally sort into high-pay firms.

When interpreting these results, it is important to keep in mind that the conditional ranks also capture differential sorting into employment. Appendix Figure A3 compares the rank-rank relationship in the full sample to the corresponding relationship for the sub-sample of children with an employer. The steepening of the rank-rank relationship below the 20th percentile is less pronounced when conditioning on the employment margin. In general, however, both series are very similar, exhibiting non-linearities at both tails. The patterns we document are therefore mostly driven by differences in which firms people work for, rather than whether they work at all.

We further assess the sensitivity of these results to different definitions of income in Appendix C.1. For instance, several children may still be pursuing graduate degrees between the age of 25 and 29. As a result, we might mischaracterize the role of firms for lifetime mobility when measuring income and employers too early. Reassuringly, using older cohorts, we find that the share of the income rank-rank slope attributable to firm sorting is largely unchanged if we measure income at older ages (29 to 33). In contrast, results are more sensitive to whether non-employment income is included or not in our measure of income. Here, we find the role of firms is even greater when non-employment income (e.g. investment income) is excluded.

4 Social Connections and Firm Sorting

4.1 Empirical Specification

Having established the magnitude of the firm sorting component, we now examine how social connections shape the distribution of workers across employers.

Econometric model. 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 social connections as predictors.[24] To do so, we endogenize realized employer-employee

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

matches and make them the dependent variable in a dyadic set-up:

$$H_{il} = \alpha_{e(i),l} + \sum_{q} \gamma^q C_{il}^q + \epsilon_{il} \tag{2}$$

where H_{il} is an indicator variable for worker i being employed by firm l, and $e(i)$ denotes the education group worker i belongs to. $\alpha_{e(i),l}$ is a match effect between education program e and firm l, and C_{il}^q is a binary variable indicating whether worker i has a social connection of type q with firm l. We include all connection types simultaneously as separate predictors. This is important since connection types are not mutually exclusive. For example, firm owners are often listed as employees of the firm they own. Firm owners also often hire more than one of their children, producing a correlation between connections to a sibling's employer and connections to a parent's business.

The coefficient γ^q reflects the impact of social connections of type q on hiring probabilities. We define social connection effects relative to a counterfactual world in which job seekers search in isolation, where workers do not use information obtained from social networks, and neither do employers. In that counterfactual scenario, workers and firms match on amenities (including geographic proximity) and observable skills, which can take on any form or number of dimensions. Relative to that counterfactual, social connections can affect the probability a worker matches with a specific employer for several reasons. First, social connections may help reveal to employers that a job applicant is a particularly good fit for the job (i.e., that they possess unobservable skills that are particularly valuable to that employer), thereby solving problems of imperfect information. This can operate via a job referral from a trusted employee, or if the firm owner knows the potential worker personally and has information about their ability that other employers do not possess. Second, social connections may help reduce search frictions in the labor market, helping disseminate information about job opportunities within social networks. Third, connected workers may benefit from preferential treatment (e.g. favoritism, nepotism). Equation (2) provides us with estimates of reduced-form social connection effects that conflate these different mechanisms.

Identification. In our context, there are at least two important threats to identification. First, children and the people they are connected to (parents, siblings, etc.) likely have similar skills, similar preferences for job attributes, and may pursue similar careers. Hence, two connected individuals may work for the same employers for reasons unrelated to the fact they know each other. Such unobserved correlations in skills would produce positive estimates of γ^q even in the absence of social connections effects. A second threat to identification is that children and their parents may look for jobs in the same local labor markets. Then, the probability of a child working for the same employer as their parent would exceed what would prevail under a pure random allocation of workers to firms at the national-level, again producing spurious positive γ^q coefficients.

We address these challenges with the inclusion of education-by-firm fixed effects $\alpha_{e(i),l}$ in equation (2), in which case the γ^q coefficients are identified by comparing "classmates" – people who hold the same degree from the same institution – to each other, as in Kramarz and Skans (2014). Intuitively, consider a comparison between two people who graduated from the same program and therefore have very similar resumes, one of which is socially connected to firm l and the other is not. The coefficient γ^q will be positive if the connected individual is more likely to work at firm l than their classmate who has the same degree but is not connected to firm l. This approach is very flexible in that it allows for any patterns of sorting across firms among sets of people with a given degree. It does not require explicit measures of workers' skills or of firms' demand for skills.

Since education groups e are based on specific post-secondary institutions, the education-by-firm fixed effects largely absorb variation in terms of which local labor market students search in: most people who studied in Quebec City look for jobs in Quebec City. Moreover, they account for patterns of intergenerational occupational following to the extent that fields of study are a key determinant of occupation sorting. This is because the inclusion of education-by-firm fixed effects implicitly restricts comparisons between people who chose to pursue the same type and level of education, and therefore have very similar human capital and career prospects. The education-by-firm fixed effects also account for specific education-to-firm pipelines that may exist independently of geographic or academic factors.[25] This includes the effect of school-based social connections that are shared equally by all students in a given program (e.g., the presence of a faculty member who has strong industry connections). As a result, these fixed effects may absorb some education-related connection effects.

[25] Oyer and Schaefer (2016) demonstrate that lawyers segregate into firms on the basis of which law school they graduated from, and Zimmerman (2019) find that graduates of elite colleges in Chile disproportionately occupy leadership positions at publicly traded firms. The proximity between universities and high-wage firms (Weinstein, 2022), as well as employers' on-campus recruiting strategies (Weinstein, 2018) also affect hiring outcomes.

4.2 Estimation

The estimation dataset for the dyadic set-up of eq. (2) is a list of all possible pairwise combinations of workers and firms.[26] 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. (2) on the full dataset is computationally impossible.

We utilize two features of the data and empirical specification to make our estimation computationally feasible. First, we can retrieve consistent estimates of γ^q on a smaller subsample of the dataset. This is because education-firm pairs with $\bar{C}^q_{e,l} = 0 \forall q$ (i.e., no one from education group e is connected to firm l) do not contribute to the identification of γ^q (Kramarz and Skans, 2014).[27] We can therefore estimate γ^q using *only* the subsample of education-firm pairs for which at least one connection of any type exists without loss of identifying variation. That is, estimates of γ^q for the full sample are exactly the same as those one obtains from estimation on the subsample in which we have discarded cells with $\bar{C}^q_{e,l} = 0 \forall q$.

Second, we partition the sample of children into broad education categories using 2-digit CIP codes (indexed by m), and estimate the model separately by category m.[28] In other words, we allow for heterogeneous effects of social connections by category m, and the estimating equation becomes

$$H_{il} = \alpha_{e(i),l} + \sum_q \gamma^q_{m(e(i))} C^q_{il} + \epsilon_{il}. \tag{3}$$

Note the education-by-firm fixed effects $\alpha_{e(i),l}$ are not indexed by m. This is because we form categories using high-level groupings of fields of study (2-digit CIP codes), so that granular education groups e are nested within categories m. For example, one such category is "Engineering", which includes any program of study at any level in some subfield of engineering. This categorization results in just over 300 categories m, and therefore we retrieve about 300 estimates of γ^q_m for each connection type.

[26] 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.

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

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

4.3 Estimates of Social Connection Effects

Figure 3 plots the distributions of estimates of the social connections coefficients γ_m^q, separately by connection type. Using the median of the distribution as a summary measure, we find the probability of working for a given employer increases by close to 4 percentage points if one's parent works there. Most coefficient estimates for connections to a parent's employer range between 0.026 (10th percentile) and 0.053 (90th percentile).[29] For comparison, the main estimate of that same parameter is 0.061 in Kramarz and Skans (2014), who use Swedish data. Given that the unconditional probabilities of matching with specific firms are very small, a 4 percentage point increase is a very large effect. Parental connections clearly matter for early career jobs.

Results in panel (b) indicate that the impact of parental connections based on ownership status is an order of magnitude larger than for connections based on employment relationships. When a parent owns a firm, the change in their child's hiring probability range between 0.279 (10th percentile) and 0.649 (90th percentile), with a median 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-based 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 their impact on hiring patterns being dramatically different.

Panels (c) and (d) present corresponding estimates for connections via siblings. Children are 2 percentage points more likely to work at a firm when one of their siblings works there, relative to people with the same education but no connection to the firm. The impact of social connections via sibling business owners (panel d) is sizable – the median of γ_m^q is 0.0654 – but quite dispersed across subgroups. There is a non-negligible share of estimates with a negative sign, indicating that in some cases siblings may avoid working for one another.

We report results for parent's past co-workers in panels (e) and (f). The impact of social connections via such social contacts is positive, but much smaller in magnitude than it is for close family members: the median estimate is 0.4 percentage point when the social contact works at a firm, and 0.001 when they own the firm. Estimates of social connection effects

[29] Roughly 95% of the estimates shown in panel (a) have p-values below 0.05 based on heteroskedasticity-robust standard errors.

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

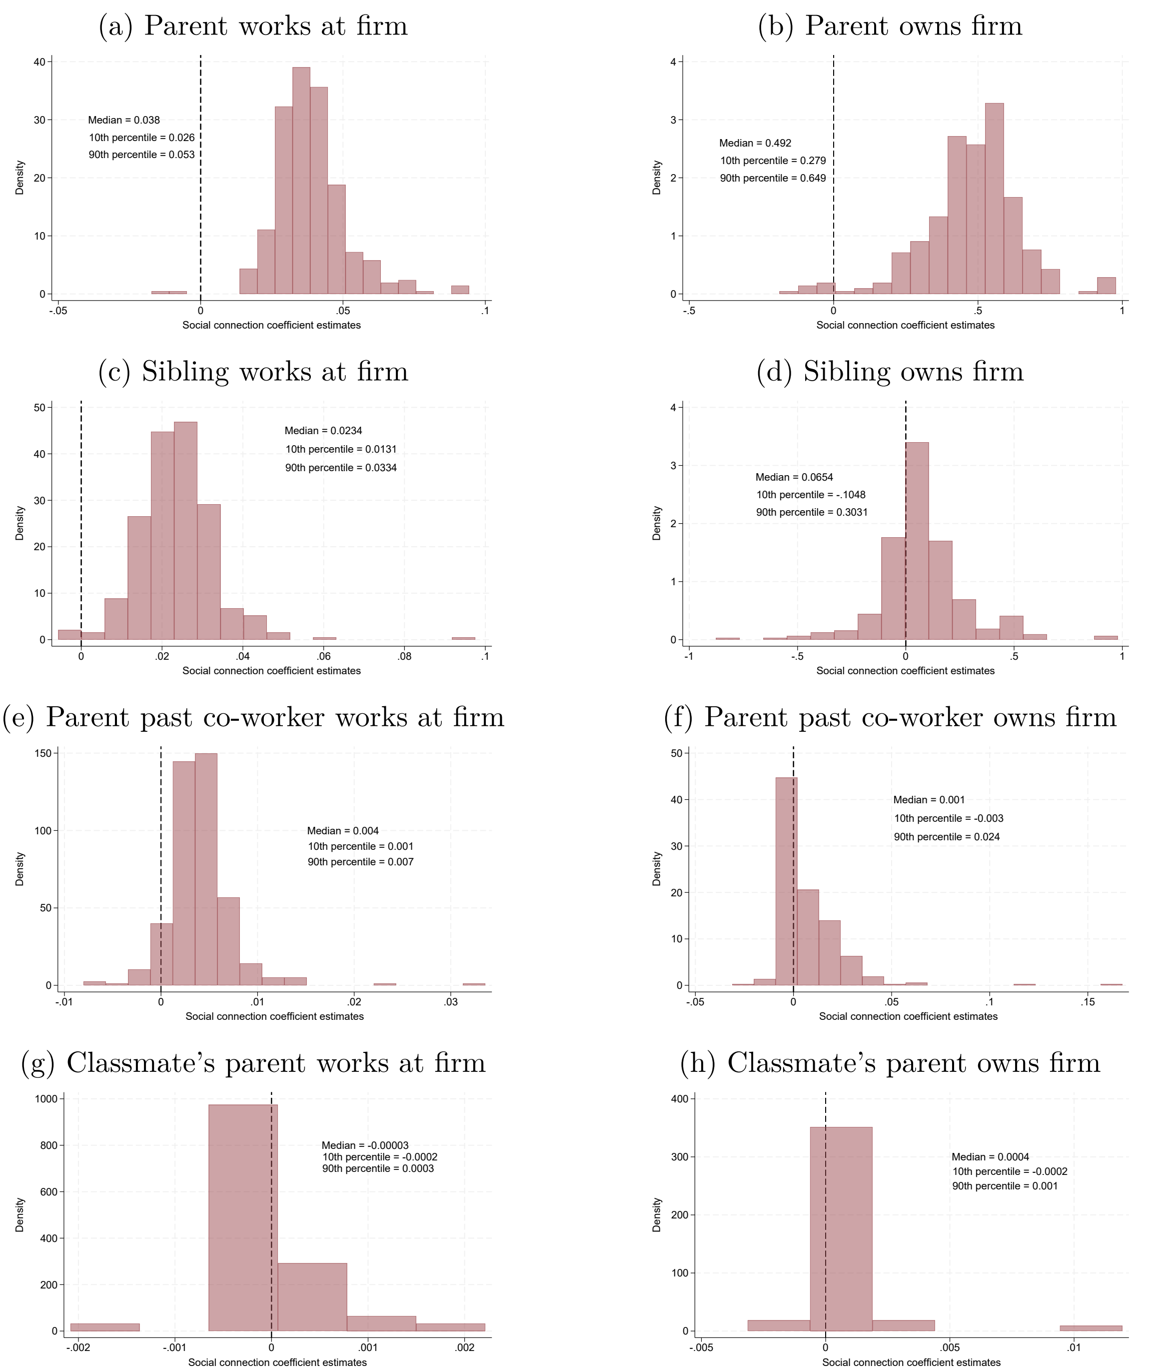

Notes: This figure presents the distribution of estimated social connection coefficients γ_m^q from equation (3). Each panel plots the distribution for a different connection type. Note the number of estimates in panels (g) and (h) is considerably smaller since these connections only apply to groups m with some post-secondary education.

via the parents' of a child's graduation cohort peers are presented in panels (g) and (h).[30] While the descriptive statistics presented in Figure 1 show it is quite common for children to share an employer with one of their classmate's parent, our results based on the dyadic framework suggest these social contacts have no impact on firm sorting. This comparison highlights the importance of controlling for skills when estimating social connection effects.

To unpack some of the heterogeneity in social connection effects, Appendix Figure A5 shows that social connections effects are larger for children without post-secondary education – except perhaps for connections to parents' past co-workers – a result also in line with the findings in Kramarz and Skans (2014). Since children from high-income families are more likely to pursue post-secondary education (see Figure A6), they may be subject to lower social connections effects on average. Figure A7 – which plots the average of $\hat{\gamma}^q_{m(i)}$ separately by connection type and parental income rank – shows that this is indeed generally the case.

4.4 Robustness

If parents and children share unobserved firm-specific skills, then the coefficients γ^q_m may reflect hiring on the basis of correlated unobserved skills rather than social connections. It might also be the case that even within narrow education groups different students look for jobs in different cities depending on where their parents live. For instance, a child who grew up in Saskatchewan and completed a Bachelor's degree in Ontario may want to move back to their home province after graduation, whereas most of their classmates would remain in Ontario. We conduct two series of empirical exercises to evaluate whether these confounds severely bias our estimates of social connection effects.

Timing of job matches and future connections. First, we focus on parental employment and allow the effect of social connections to differ by the length of the parent's tenure at the job, *including negative tenure*, in the estimating equation.[31] That is, we include future connections (i.e. the parent has yet to start working at firm l) as placebo connections. To do so, we re-estimate the dyadic set-up of eq. (2) 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 val-

[30] The distributions are less dispersed in these two cases because these connection types only exists for children who graduated from a post-secondary program, and therefore only exist for categories m corresponding to some post-secondary education.

[31] For this exercise, we do not consider connections based on ownership since a child cannot possibly start working at a firm before it exists, that is before the business owner starts it.

ues of job tenure at that point in time. If the effect does operate via social connections, then we would expect a sharp break in a child's hiring probabilities around the time their parent joins the firm. The coefficients on future connections capture a combination of correlated unobserved skills, locational preferences, and unobserved shared connections. For example, future parental connections may 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. It may also be the case that the child is the person providing a referral for their parents.

As before, we estimate the model separately by broad educational groups m. Results are summarized in Figure 4.[32] For visual clarity, for each value of job tenure, we plot the median of coefficient estimate across subgroups using blue lines, and use bars to present the range between the 10th and 90th percentiles of the distribution of coefficient estimates. Positive values of job tenure T mean that a parent of child i was working 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 tenure 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.

We find the impact of parental connections is larger the longer the parent has been at the firm. When the parent has been at the firm for 5 years or more, the median change in hiring probabilities is close to 0.08, almost twice the size of our main estimates pooling all tenure lengths together. Estimates are slightly lower in magnitude for tenure of 4, 3, 2, and 1 year, on average, and are more dispersed, as we may expect given the fact there are far fewer of these short-tenure connections than long-tenure ones when a child is 25 years old. Then, there is a sharp drop in the magnitude of the estimated coefficients between tenures of 1 and -1 years.[33] The distributions of coefficients for future connections are more concentrated, and gradually shift toward zero as we move away from the present, similar to the findings in Staiger (2025). To the extent correlated unobserved skills and locational preferences are time-invariant, then they should affect all coefficients the same way. In contrast, the set of firms one has access to via broader unobserved social networks likely changes over time. Putting these two arguments together suggests the slow decline of coefficient sizes between job tenure of -1 to -5 is mostly due to changing social connections. Most estimates for tenure of -5 are very close to zero. We take this as evidence that any bias in our social

[32]The complete set of distributions of estimates is presented in Appendix Figure A8.

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

Figure 4: Distributions of Social Connections Effects on Hiring Probabilities, by Parent's Tenure At Employer

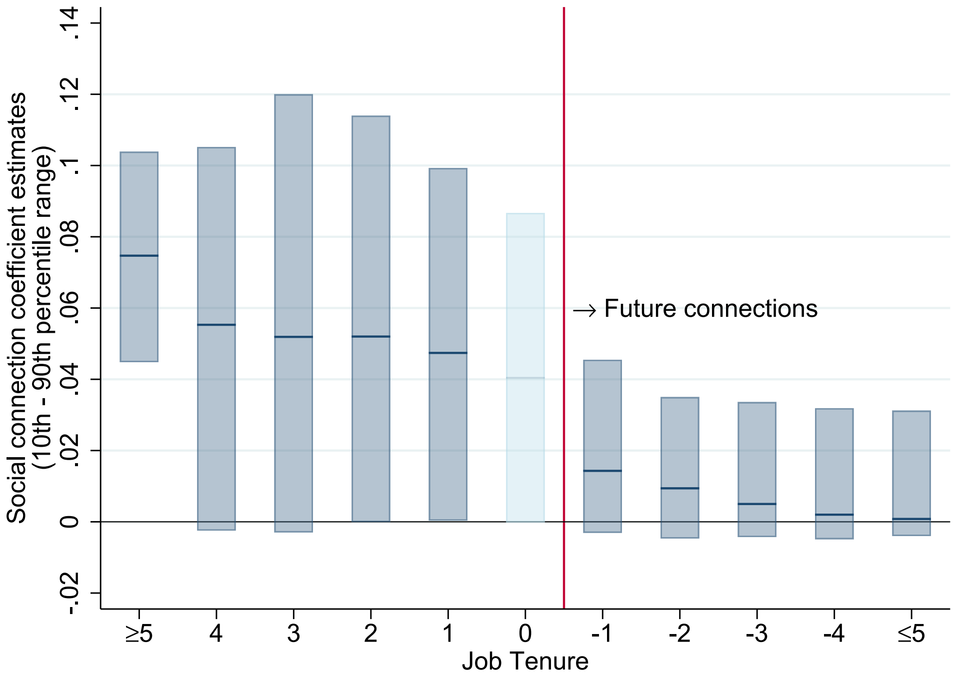

Notes: This figures plots key moments of the distribution of social connection coefficient estimates for different values of job tenure. The median is depicted by thin blue lines, and vertical bars indicate the 10th to 90th percentile range. he estimating equation is $H_{il}^{25} = \alpha_{e(i),l} + \sum_T \gamma_{m(e(i))}^T C_{il}^T + \epsilon_{il}$, where H_{il}^{25} is an indicator for worker i working at firm l when they are aged 25, and C_{il}^T indicates whether a parent has a job tenure of T years at firm l the year their child is aged 25.

connection estimates is likely minimal.

Within-neighborhood permutations. Second, while the education-by-firm fixed effects absorb variation in geographic preferences to a large extent, it remains possible that our estimates of social connection effects are biased due to differences in terms of which local labor market graduates search in. To evaluate this possibility, we run a placebo-type analysis in which we permute each child's true employer (i.e., the value of H_{il}) with one of their childhood neighbor's employer. That is, for each child in our sample we randomly select a "control" child that grew up in the same 6-digit postal code – which amounts to living on the same block-face or even in the same apartment building in denser areas – and use that control child's employer as the outcome in eq. (3). Candidate control children exclude siblings, but include children from birth cohorts outside of our main analytical sample.[34] We present the results of this analysis in Appendix Figure A9. For all eight connection types, the distributions of estimates are all tightly clustered around zero.

4.5 Quantifying the Importance of Social Connections

Method. What is the economic significance of social connection effects for income mobility? To answer this question, we collapse equation (3) at the firm-by-parental income rank level:

$$s_{l|p} = E[H_{il}|l,p] = \underbrace{E[\alpha_{e(i),l}|l,p] + \sum_q E[\gamma^q_{m(e(i))} C^q_{il}|l,p]}_{\tilde{s}_{l|p}} + E[\epsilon_{il}|l,p]. \tag{4}$$

where $\tilde{s}_{l|p}$ is a counterfactual (predicted) distribution of children across firms based on predictors included in the model. In order words, we split the observed distributions of children across firms by parental income into a part that can be explained by differences in observables (education and social connections), and a part that cannot. Equation (4) also makes clear that the quantitative importance of connections for the firm sorting component of income mobility depends on the effect size (the magnitude of γ^q_m), but also on the incidence of connections (the values of C_{il}).

The associated counterfactual firm sorting component is obtained by substituting the

[34]Since education is unobserved for children in earlier birth cohorts, permuting the right-hand side variables is infeasible.

counterfactual distributions $\tilde{s}_{l|p}$ for the real ones:

$$\tilde{\Delta}_p = \sum_l \delta_l \tilde{s}_{l|p}. \tag{5}$$

Calculating all elements of $\tilde{s}_{l|p}$ requires estimates of $\alpha_{e,l}$ for all possible education-firm pairs, including those not in the estimation subsample. Appendix B.4 describes how we can back-out all values of $\alpha_{e,l}$ using estimates of γ_m^q as well as some other moments of the data.

Results Figure 5 presents a breakdown of the firm sorting component based on selection-adjusted firm effects by stacking all of the terms in $\tilde{\Delta}_p$.[35] We also tack on the residual term $r_p = \Delta_p - \tilde{\Delta}_p$ to examine how the eight types of social connections collectively account for overall patterns of firm sorting. Note the model in equation (3) does not include information on parental income. That is, we predict individual assignment to firms using information on social connections and education, and then examine the extent to which the counterfactual sorting component $\tilde{\Delta}_p$ can reproduce the true firm sorting component Δ_p. To summarily 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 (2025).[36] This allows to take into account any non-linearity in their explanatory power across the parental income distribution.

We find that the eight types of social connections we consider collectively explain just over 13% of the firm sorting component. For benchmarking purposes, we also plot the share attributable to differences in education in panel (a). Education accounts for 45% of the firm sorting component, indicating the overall explanatory of social connections is about a third that of detailed education. One important caveat is that the share attributable to education may not just capture a causal effect of education, by could also reflect the influence of unobserved heterogeneity in academic ability, as well as other potential correlates of educational decisions. Therefore, a comparison of these two components likely provides a lower bound on how social connection effects compare to education effects. The fact we only observe a subset of all possible social connections further reinforces that point. Finally,

[35] Corresponding results for unadjusted firm fixed effects are presented in Appendix Figure A10. Results are very similar.

[36] 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 $\sum_l \delta_l E[\alpha_{e(i),l}|l,p]$ 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 ensures that all shares sum to 1.

Figure 5: Decomposition of the firm sorting component

(a) Pooling all connection types

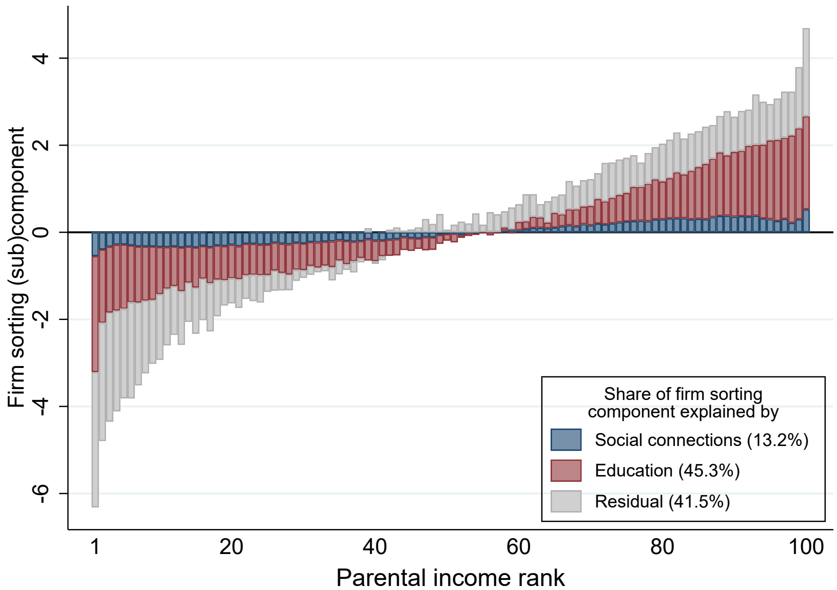

(b) Separately by connection type

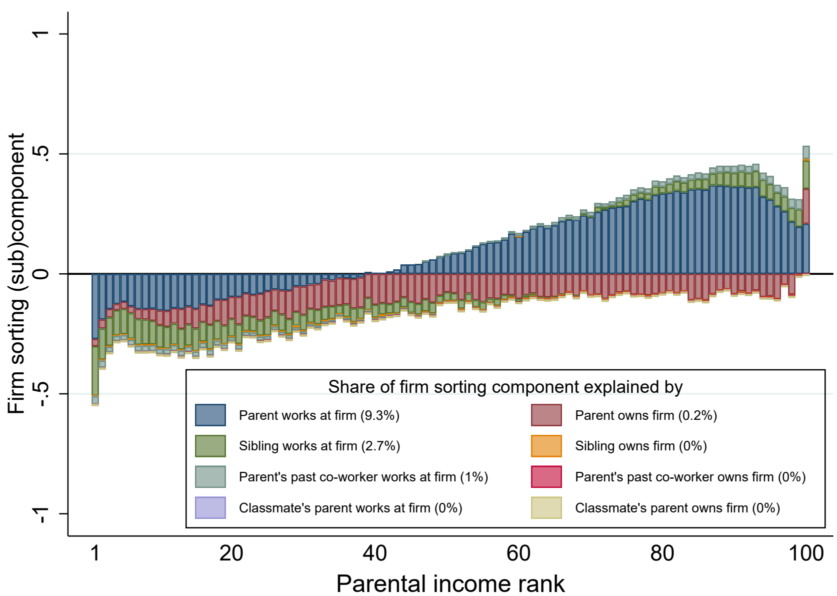

Notes: This figures plots counterfactual firm sorting components based on equations (4) and (5). Panel (a) shows components attributable to education ($\sum_l \delta_l E[\alpha_{e(i),l}|l,p]$) and to social connections pooling all eight types ($\sum_l \delta_l \sum_q E[\gamma^q_{m(e(i))} C^q_{il}|l,p]$), as well as residual unexplained variation in firm sorting. Panel (b) decompose the social connection component, breaking down the contribution of each connection type. Values of δ_l are selection-adjusted firm effects.

the residual component accounts for almost half of the firm sorting component. Factors we do not measure that are part of this residual include, for example, within-education group differences in academic ability, unobserved social connections, as well as differences in preferences for job attributes across parental income groups.

In panel (b), we break down the social connection component into parts attributable to each connection type. We find the lion share of social connections' explanatory power comes from connections to a parent's employer (9 percent points out of 13). Not only is this type of connections very common, but the quality of firms children have access to this way is strongly correlated with parental income.[37] In contrast, connections to a sibling's employer explain only 3% of the firm sorting component. Connections to firm owners – either a parent, a sibling or a classmate's parent – have virtually no explanatory power in total. Between the 45th and 99th percentiles of the parental income distribution, if anything, connections to firms parents own have *negative* explanatory power. That is, on the basis of such connections, we'd predict children in these parental income groups would have below-average firm-sorting components, whereas in reality their firm-sorting component is above-average. This is because most family firms are low-pay firms. In contrast, the contribution of connections to family firms is positive and economically significant for the top 1 percentile of parental income.

The values presented above likely constitute a lower bound on the role of social connections. First, while our measures of connections are relatively broad, there exists many connection types (e.g., friends from high school) we cannot observe and measure. Second, it might be that some of the variation we attribute to differences in education operates via the use of social networks built in school, rather than to differences in human capital.[38] In this case, some fraction of the education component should rather be interpreted as a social connection component. We consider this possibility in Appendix D, where we decompose the education-by-firm coefficients into an institution-by-firm fixed effect (capturing factors due to geographic proximity between a school and a firm, as well as reputation effects), a field of study-by-firm fixed effect (capturing each firm's propensity to hire graduates from a given field), and an idiosyncratic match effect between program e and firm l. The first two components are unlikely to capture social connection effects since they reflect attributes shared by many students who certainly do not know each other – people from the same

[37]This correlation is not completely mechanical since parental income is measured when the child is aged 15-19, whereas parents' employers are measured when the child is 25-29.

[38]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 generally only based on grades.

university but in different programs, and people in the same field of study but from different places in the country. More plausible is that social connections between classmates operate through residual match effects. We find this residual variation accounts for a non-negligible 5% of the firm-sorting component.

5 On the Purpose of Social Connections in Hiring

Why are workers more likely to match with a firm they are socially connected to? There exists several classes of explanations, each of which have different implications for the effect of connections on the earnings of connected workers and on the performance of firms who hire these connected workers. The mechanisms through which connections affect hiring probabilities also likely differ by connections types. For instance, in some cases a worker may receive an explicit job referral from someone currently employed at the firm.[39] But in cases where workers are connected to the firm owner, referrals are unlikely to be relevant as the contact person may be the one making hiring and pay decisions. Below, we conduct two sets of empirical analyses to shed light on potential mechanisms. Specifically, we examine whether connected workers earn more, less or the same as similar unconnected colleagues, and whether firms that hire connected workers perform better, worse, or the same as those that hire similar unconnected workers.

5.1 Do Connected Workers Earn More Than Their Colleagues?

Estimation. We begin our empirical analysis to distinguish between alternative mechanism by looking at how social connections affect workers' compensation relative to their co-workers. We do so using the following regression equation:

$$y_i = \sum_q \theta^q C^q_{ij(i)} + \omega_{e(i)} + \kappa_{j(i)} + \varepsilon_i \tag{6}$$

where y_i is worker i's income rank, $C^q_{ij(i)}$ indicates whether worker i has a social connection of type q with the firm $j(i)$ that hires them (i.e., their modal employer between the age of 25-29), and $\omega_{e(i)}$ is a set of education group fixed effects. The coefficients θ^q capture the income advantage that accrues to workers with a social connection of type q relative to unconnected

[39]Brown et al. (2016) provides a detailed list of empirical predictions associated with different models of employment referrals.

worker. Although estimates of θ^q do not have a causal interpretation, they are net of selection on very detailed measures of education. We report estimates both from models with and without firm fixed effects to get a broad view of where any income advantage associated with social connections may come from.[40] The inclusion of firm fixed effects insures that parameters θ^q and $\omega_{e(i)}$ are estimated using within-firm variation, and therefore are based on comparisons among colleagues. When firm fixed effects are omitted, the coefficients capture both within-firm income advantages as well as income gains arising from access to better paying firms.

Results. We report estimates of the social connection coefficients θ^q in Figure 6. When firm fixed effects are omitted, children who work for the same employer as their social contact (either a parent, a sibling, a parent's former co-worker, or a classmate's parent) have income levels substantially higher – between 5 and 7 percentile ranks – than unconnected workers with the same degree. The corresponding within-firm income advantages (i.e., when accounting for firm fixed effects) are considerably smaller, closer to 1 or 2 percentile ranks. This means that the income advantage associated with working at a contact's employer operates mostly through access to better-paying firms, a conclusion that Staiger (2025) also reaches. Strikingly, patterns go the opposite direction for connections to a firm's owner. For instance, children working for their parents have similar income as unconnected workers (at other firms) with similar education, but much higher incomes (about 15 percentiles higher) than their colleagues. The fact that the coefficient from a model 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 firms offer higher remuneration to socially connected workers relative to other same-aged employees?[41] One possibility is that connected workers are more productive. Such productivity differentials may arise for several reasons: social connections may help reveal private information about unobserved skills at the time of hiring, connected workers may have better access to information about how to navigate

[40]When firm fixed effects are included, singleton firms are necessarily dropped. For robustness, in Figure A11 we report results in which singleton firms are kept in the estimation sample by pooling them into ventiles of the distribution of AKM firm effect. See Appendix B.1 for details of how firms are grouped.

[41]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 likely only receive employment income. In Appendix Figure A12, we re-estimate the model but using employment income alone as the outcome. The coefficients for parent-owned and sibling-owned businesses, not accounting for firm fixed effects, are *negative* when non-employment income is excluded.

Figure 6: Social Connections and Income Ranks

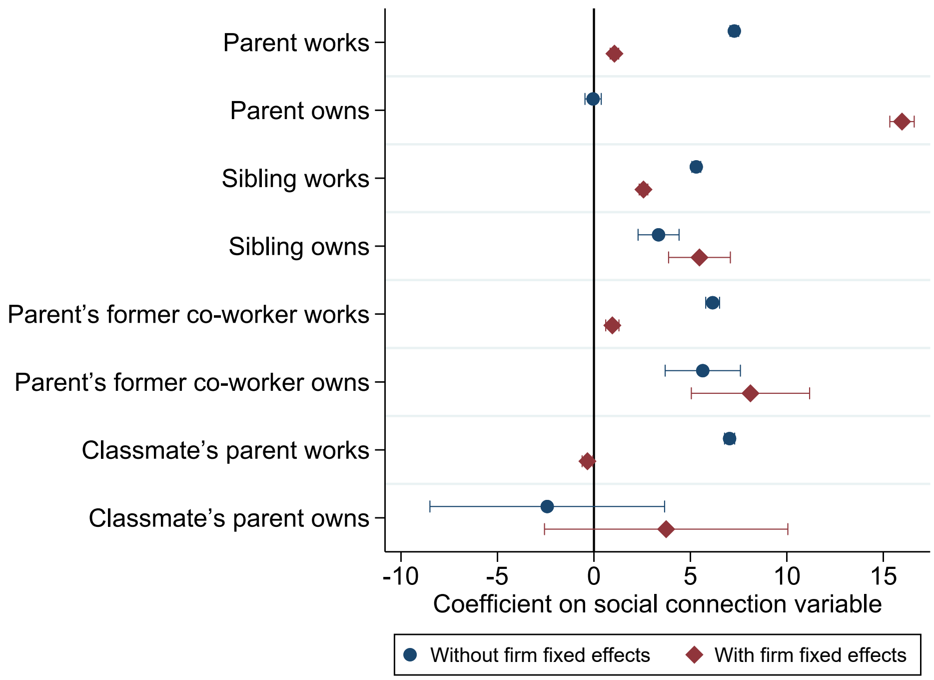

Notes: This figure plots coefficient estimates from equation (6), separately by connection type and by specification that do or do not include firm fixed effects. All models include education fixed effects. Whiskers indicate 95% confidence intervals based on heteroskedasticity-robust standard errors.

internal labor markets, or workers may be productive when working with people they know. Alternatively, connected workers may benefit from preferential treatment (e.g., favoritism). Ideally, one would directly measure worker productivity to distinguish between these competing explanation. Without such data, however, we flip the perspective and examine whether firms gain from hiring connected workers.

5.2 Social Connections and Firm Performance

How do social connections affect firm performance? An important challenge is that worker sorting to firms is not random and may depend on unobserved determinants of performance. We adapt our classmates research design that exploits highly granular data on education to estimate the effects of hiring connected workers on firm performance. Our approach is to form comparison sets g in which all firms – both "treatment" and "control" firms – hired students of the same education program e in the same year h, but in some cases the worker had a social connection with the employer, and in other cases they did not. This narrows the comparison between firms who likely faced similar market conditions since they chose to hire workers with very similar skills in the same year.

Estimation. Specifically, we estimate the equation:

$$v_{jgt} = \sum_q \sum_\tau \beta_\tau^q \left(\mathbb{1}\{t - h_g = \tau\} \times C_{jg}^q\right) + \phi_{jg} + \phi_{gt} + \epsilon_{jgt} \tag{7}$$

where v_{jgt} is an outcome for firm j in comparison set g in calendar year t. Note that firms can appear in multiple comparison sets if they hire different workers in different years. Then, C_{jg}^q indicates whether firm j hired a worker with social connection type q in comparison set g, and h_g is the year all firms in comparison set g hired at least one worker from education group e. ϕ_{jg} are firm-by-comparison set fixed effects, and ϕ_{gt} are comparison set-by-year fixed effects. The latter account for any time trend that is shared by firms that hired workers from the same program in the same year.

The coefficients of interest are β_τ^q, which are on interactions between dummies for time relative to the event year h_g (omitting the dummy for $\tau = -1$), and treatment status C_{jg}^q. As in our previous analyses, all eight connection types are simultaneously included in the model. We 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), as recommended in Schmidheiny and Siegloch (2023) to be able to estimate the relevant effect dynamics.

The main outcomes we track are log revenue per worker and log value-added per worker, which we winsorize below the 1st percentile and above the 99th percentile of their distributions to reduce the influence of outliers. For internal consistency, we focus on hiring events when a worker from our main sample of children born in 1987-89 is joining their modal employer. Since an 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.

Results. Results are presented in Figure 7. Firm performance is slightly reduced after a worker joins their parent's employer, but declines precipitously in cases where a worker joins a firm their parents own. One possibility is that the child takes over the family firm and goes through a transition phase as he or she learns how to manage the company. The patterns we observe, however, are not consistent with the change being transitory: 5 years after the child is hired log revenue per worker remains much lower than at $\tau = -1$. There is no systematic change in per worker firm outcomes for the remaining connection types.[42] Overall, there is no connection type for which firm performance increases when a connected worker is hired.

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.[43] This effect could capture private benefits of control in family firms that manifest over much longer time horizons (Di Porto et al., 2024). Still, an important caveat is in order: all of these findings only apply to small firms and may not generalize to larger firms.

[42] Our main outcomes of interest are on a per worker basis. In Appendix Figure A13 we examine trends in firm scale by presenting results for log revenue and log payroll outcomes. We see that disentangling causal effects of hiring connected workers from pre-existing trends in firm scale is very challenging. For example, firms hiring the owner's children were contracting before they hired the connected worker. After the worker joins the firm, revenue continue to decrease relative to control firms, whereas payroll increases. In contrast, for connections to a sibling's employer, both revenue and payroll were increasing prior to the event (indicating firms were expanding), and continue to do so in similar proportions after a connected child is hired.

[43] Firm owners may also hire their child for tax avoidance purposes. However, most ways to do so require the child to be living with the parent. Since we focus on modal employers at 25-29, we believe our findings aren't driven by tax incentives.

Figure 7: Firm Performance and Hiring of Connected Workers

(a) Connections via parents

(b) Connections via siblings

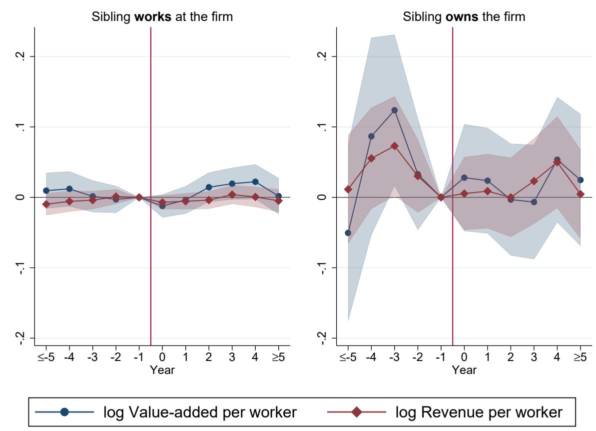

(c) Connections via parents' past co-workers

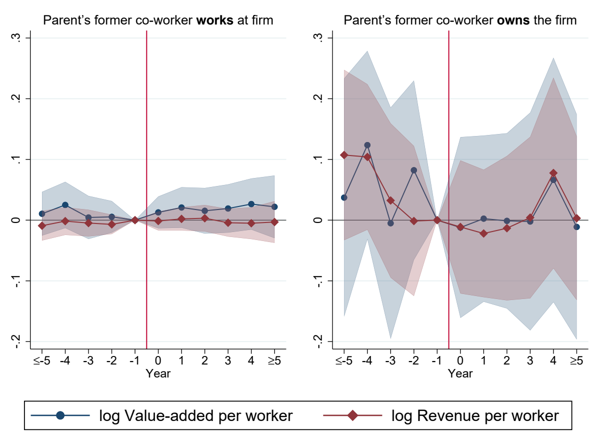

(d) Connections via classmates' parents

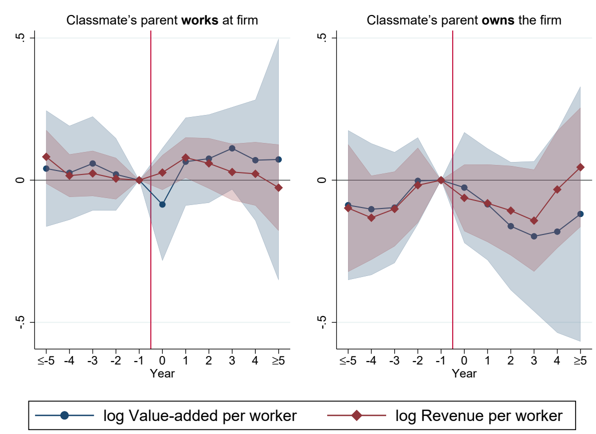

Notes: The figure presents event-study estimates of changes in firm performance around the time a connected workers is hired. Point estimates are based on equation (7). Shaded areas indicate 95% confidence intervals based on standard errors clustered by comparison set g.

6 Discussion and Conclusion

Using a detailed matched employee-employer-parent-child database for Canada we document five stylized facts. First, which firms workers match with matters for intergenerational income mobility. Differential representation at high-pay employers explains roughly half of the transmission of income across generations, as measured by the income rank-rank relationship.

Second, using a classmates research design, we find that social connections significantly impact job match probabilities. These effects differ considerably by connection type. The likelihood of working at a given firm increases by about 3 to 5 percentage points when one's parent works there. The effect size is larger by an order of magnitude when a parent is the owner of the firm. Social connection effects are smaller for sibling and non-family connections.

Third, connections to family firms account for almost none of the firm sorting component of the rank-rank relationship, except at the very top of the income distribution. Connections to a parents employer are quantitatively more important, explaining close to 10% of the firm sorting component. Our results highlight the importance of differing connection incidence across the parental income distribution in explaining the persistence of income across generations.

Fourth, workers who have a social connection to their employer earn more, on average, than their coworkers who do not have such a connection. Still, for connections to a social contact's employer, most of the income advantage comes from access to better-paying employers, not access to better-paid positions within the firm.

Finally, that social connections contribute to intergenerational (im)mobility raises questions over their efficiency. We show that firm revenue per worker and value-added per worker decline after a connected child joins a firm their parents either own or work at, and persist for at least a 5-years. These results indicate that workers connected though their parents receive preferential treatment rather than making the connected firm more productive.

While much work on intergenerational mobility has focused on human capital, our results demonstrate an important role for labor market frictions. They imply the frictions a worker faces differ by parental background. Some limitations of our study suggest important directions for future work. One could be incorporating heterogeneous returns to job search effort that differ by parental income into a search and matching model following (Faberman et al., 2022). Another avenue would be to incorporate labor market frictions in the next generation of models of intergenerational mobility. There remains much to learn about whether and how differences in labor market outcomes across workers are driven by differences in the labor market frictions they face.

References

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

Agostinelli, Francesco and Matthew Wiswall, "Estimating the Technology of Children's Skill Formation," *Journal of Political Economy*, 2025, *133* (3), 846–887.

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.

— and — , "Human Capital and the Rise and Fall of Families," *Journal of Labor Economics*, 1986, *4* (3, Part 2), S1–S39.

Bennedsen, Morten, Kasper Meisner Nielsen, Perez-Gonzalez, and Daniel Wolfenzon, "Inside the Family Firm: The Role of Families in Succession Decisions and Performance," *Quarterly Journal of Economics*, 2007, pp. 647–691.

Bertrand, Marianne and Antoinette Schoar, "The Role of Family in Family Firms," *Journal of Economic Perspectives*, 2006, *20* (2), 73–96.

Björklund, Anders, Mikael Lindahl, and Erik Plug, "The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data," *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," *Scandinavian Journal of Economics*, 2017, *119* (1), 72–101.

Brown, Meta, Elizabeth Setren, and Giorgio Topa, "Do Informal Referrals Lead to Netter Matches? Evidence from a Firm's Employee Referral System," *Journal of Labor Economics*, 2016, *34* (1), 161–209.

Burkart, Mike, Fausto Panunzi, and Andrei Shleifer, "Family Firms," *The Journal of Finance*, 2003, *58* (5), 2167–2201.

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The Value of Hiring Through Employee Referrals," *Quarterly Journal of Economics*, 2015, *130* (2), 805–839.

Calvo-Armengol, Antoni and Matthew O. Jackson, "The Effects of Social Networks on Employment and Inequality," *American Economic Review*, 2004, *94* (3), 426–454.

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," *Quarterly Journal of Economics*, 2016, *131* (2), 633–686.

— , — , Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

— , Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 2013, *128* (3), 967–1015.

Caucutt, Elizabeth M, Lance Lochner, Joseph Mullins, and Youngmin Park, "Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments," Technical Report, National Bureau of Economic Research 2020.

Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan, "Income Segregation and Intergenerational Mobility across Colleges in the United States," *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, Drew Johnston, Martin Koenen, Eduardo Laguna-Muggenburg, Florian Mudekereza, Tom Rutter, Nicolaj Thor, Wilbur Townsend, Ruby Zhang, Mike Bailey, Pablo Barbera, Monica Bhole, and Nils Wernerfelt, "Social Capital I: Measurement and Associations with Economic Mobility," *Nature*, 2022, *608* (7921), 108–121.

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

Childs, Stephen E, Ross Finnie, and Felice Martinello, "Postsecondary Student Persistence and Pathways: Evidence from the YITS-A in Canada," *Research in Higher Education*, 2017, *58*, 270–294.

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 — , "The Great Mobility Puzzle Resolved," *Canadian Public Policy*, 2026, *52* (1), 1–21.

— , — , 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," 2025.

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," *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.

Engzell, Per and Nathan Wilmers, "Firms and the intergenerational transmission of labor market advantage," *American Journal of Sociology*, 2025, *131* (2), 322–370.

Faberman, R. Jason, Andreas I. Mueller, Aysegul Sahin, and Giorgio Topa, "Job Search Behavior Among the Employed and Non-Employed," *Econometrica*, 2022, pp. 1943–1779.

Fischer, Alexander, Andrei Gorshkov, Tróndur M Sandoy, and Jeanette Walldorf, "Peers, Heirs and Careers: Labor Market Effects of Alumni Networks," 2025.

Forsberg, Erika, Martin Nybom, and Jan Stuhler, "Labor Market Drivers of Intergenerational Earnings Persistence," 2026.

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.

Gendron-Carrier, Nicolas, "Prior Work Experience and Entrepreneurship: The Careers of Young Entrepreneurs," 2023.

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.

Granovetter, Mark, *Getting a Job: A Study of Contacts and Careers*, University of Chicago Press, 1995.

Haeck, Catherine and Jean-William Laliberté, "Careers and Intergenerational Income Mobility," *American Economic Journal: Applied Economics*, 2025, *17* (1), 431–458.

Kaila, Martti, Emily Nix, and Krista Riukula, "Impact of an Early-Career Shock on Intergenerational Mobility," *Journal of Labor Economics*, 2025, *43* (4), 1035–1062.

Kenedi, Gustave and Louis Sirugue, "Intergenerational Income Mobility in France: A Comparative and Geographic Analysis," *Journal of Public Economics*, 2023, *226*, 104974.

Kirkebøen, Lars, Edwin Leuven, and Magne Mogstad, "College as a marriage market," Technical Report, National Bureau of Economic Research 2021.

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.

Lester, Benjamin, David A Rivers, and Giorgio Topa, "The Heterogeneous Impact of Referrals on Labor Market Outcomes," 2023.

Oreopoulos, Philip and Uros Petronijevic, "Making College Worth It: A Review of the Returns to Higher Education," *The Future of Children*, 2013, *23* (1), 41–65.

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

Pallais, Amanda and Emily Glassberg Sands, "Why the Referential Treatment? Evidence from Field Experiments on Referrals," *Journal of Political Economy*, 2016, *124* (6), 1793–1828.

Porto, Edoardo Di, Marco Pagano, Vincenzo Pezone, Raffaele Saggio, and Fabiano Schivardi, "Careers and Wages in Family Firms: Evidence from Matched Employer-Employee Data," Technical Report, National Bureau of Economic Research 2024.

Sacerdote, Bruce, "How Large are the Effects from Changes in Family Environment? A Study of Korean American Adoptees," *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.

Schmidheiny, Kurt and Sebastian Siegloch, "On Event Studies and Distributed-Lags in Two-Way Fixed Effects Models: Identification, Equivalence, and Generalization," *Journal of Applied Econometrics*, 2023, *38* (5), 695–713.

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," *Quarterly Journal of Economics*, 2019, *134* (1), 1–50.

Staiger, Matthew, "The Intergenerational Transmission of Employers and the Earnings of Young Workers," 2025.

Strobel, Stephenson, David Kanter-Eivin, Helena Son, and Adam Steiner, "The Effects of Allowing Professional Incorporation on Physician Labour Supply," *medRxiv*, 2025, pp. 2025–03.

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

Topa, Giorgio, "Social and Spatial Networks in Labour Markets," *Oxford Review of Economic Policy*, 2019, p. 722–745.

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

“Social Connections and the Persistence of Income Across Generations”

Jean-William Laliberté and Alexander Whalley

Appendix Figures

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

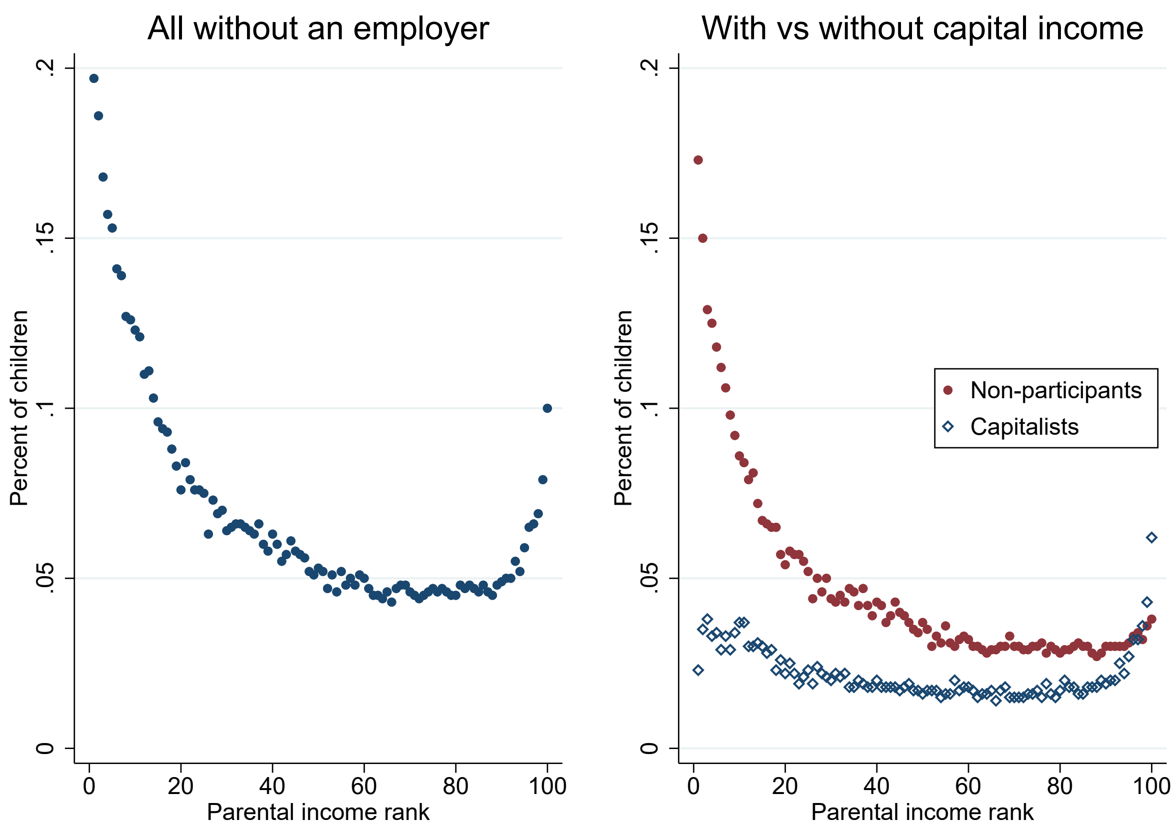

Notes: This figure shows the share of children who do not have any employer while being between the age of 25 and 29, separately by parental income rank. In the right panel, the share is split between two mutually exclusive categories: those without an employer and with no capital income, and those without an employer but with some positive amount of capital income.

Figure A2: Fraction of Children in Lone-Parent Families by Parental Income Rank

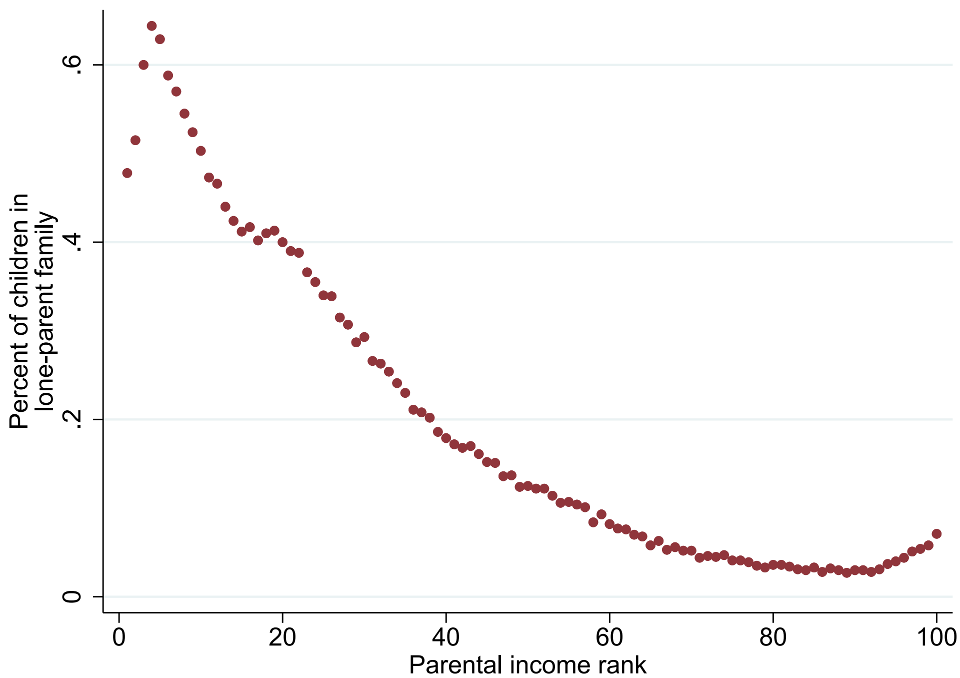

Notes: This figure presents the share of children who are in lone-parent families (i.e., have only one parent present between the ages of 15 and 19), separately by parental income percentiles.

Figure A3: Income rank-rank relationship, with and without non-employed individuals

Panel A: Unconditional rank-rank relationship

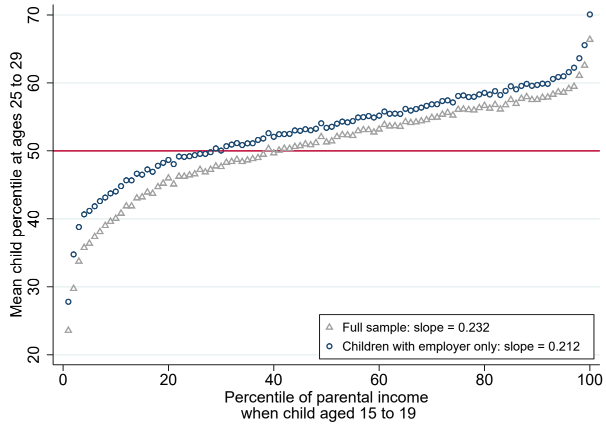

Panel B: Rank-rank relationship conditional on employers (unadjusted firm effects)

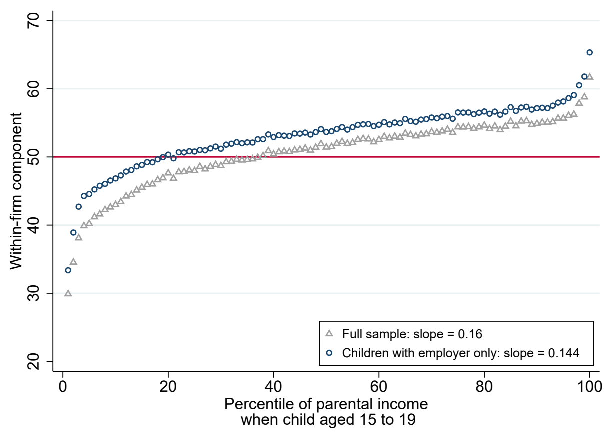

Panel C: Firm sorting component (unadjusted firm effects)

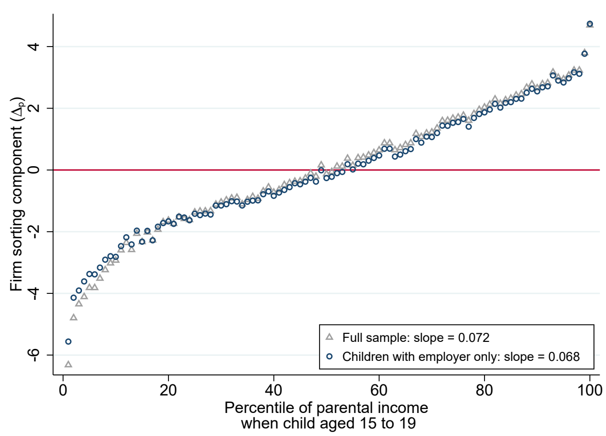

Notes: Panel (a) shows mean child income percentiles for each parental income decile. Grey triangles show unconditional means for the full sample of children born in 1987-89, whereas blue circles show unconditional means for a subsample that excludes children without a modal employer. Panel (b) shows conditional means accounting for differences in employers (using selection-adjusted firm effects), as per equation (1). Panel (c) shows the corresponding firm sorting component, also based on selection-adjusted firm effects.

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

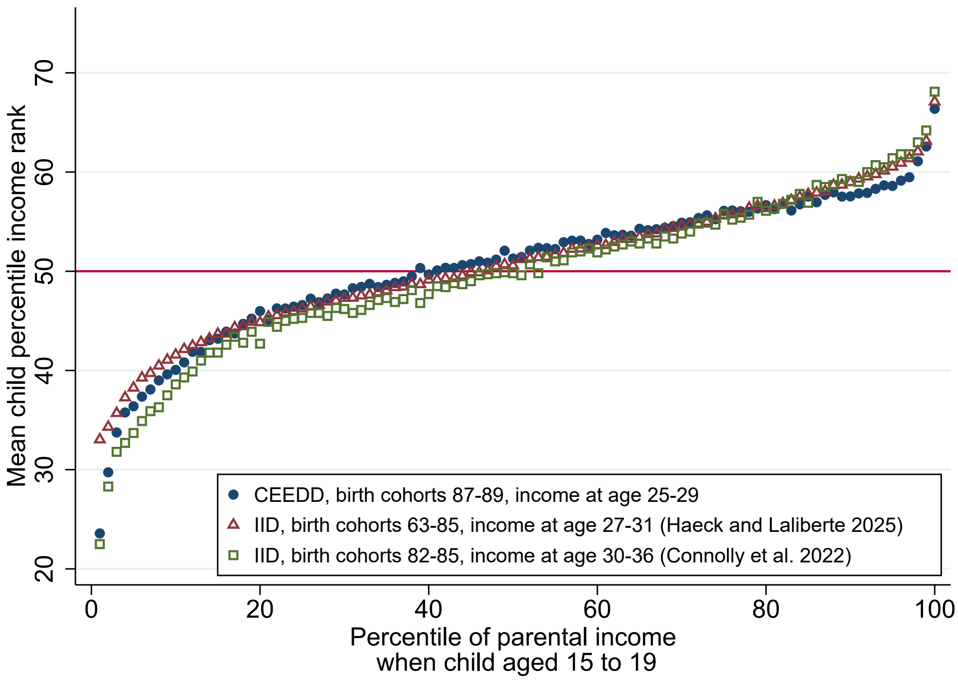

Notes: This figure shows mean child percentile income rank for three different sample. The series in blue circles corresponds to our main sample, drawn from the CEEDD. It includes children born in 1987-89, and their income is calculated when they are aged 25-29. The other two series are based on data from the Intergenerational Income Database (IID). The series in red triangles is from Haeck and Laliberté (2025). 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: Distributions of Estimates of Social Connections Effects on Hiring Probabilities, With and Without Post-secondary Education

(a) Parent works at firm

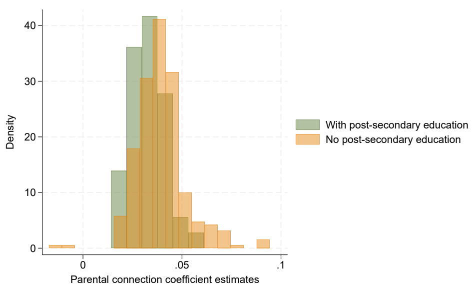

(b) Parent owns firm

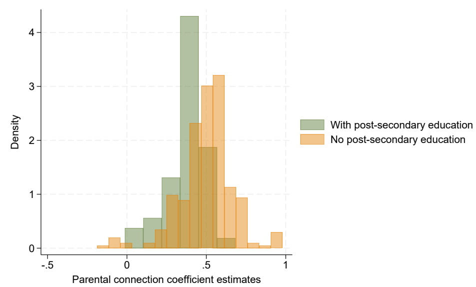

(c) Sibling works at firm

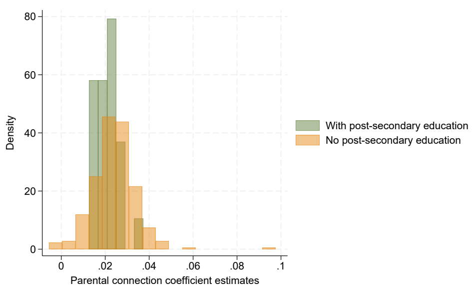

(d) Sibling owns firm

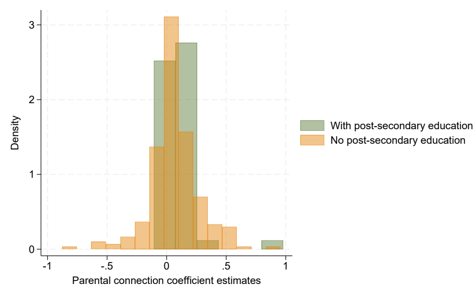

(e) Parent past co-worker works at firm

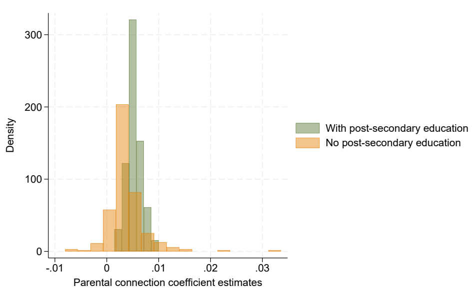

(f) Parent past co-worker owns firm

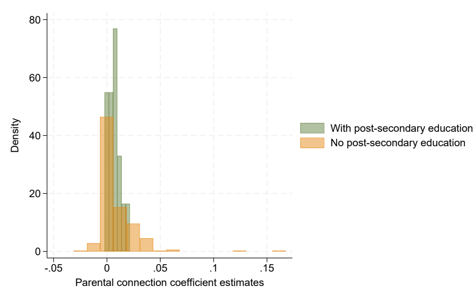

(g) Classmate's parent works at firm

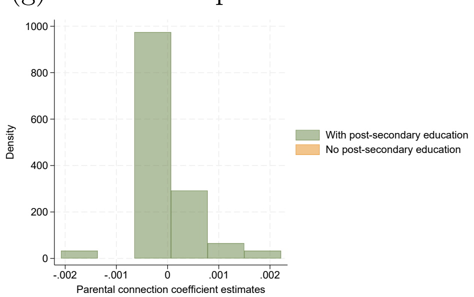

(h) Classmate's parent owns firm

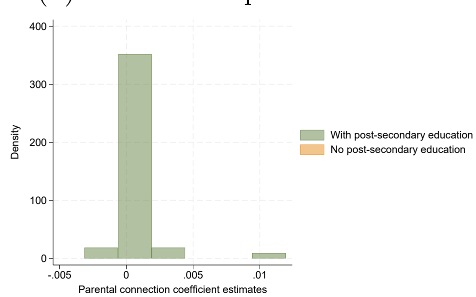

Notes: This figure presents the distribution of estimated social connection coefficients γ_m^q from equation (3), separately for subgroups m that correspond to education categories with and without post-secondary education. Each panel plots the distribution for a different connection type. Note that by definition connections via classmates' parents do not exist for children who have no post-secondary education.

Figure A6: Fraction of Children with Post-Secondary Education by Parental Income Rank

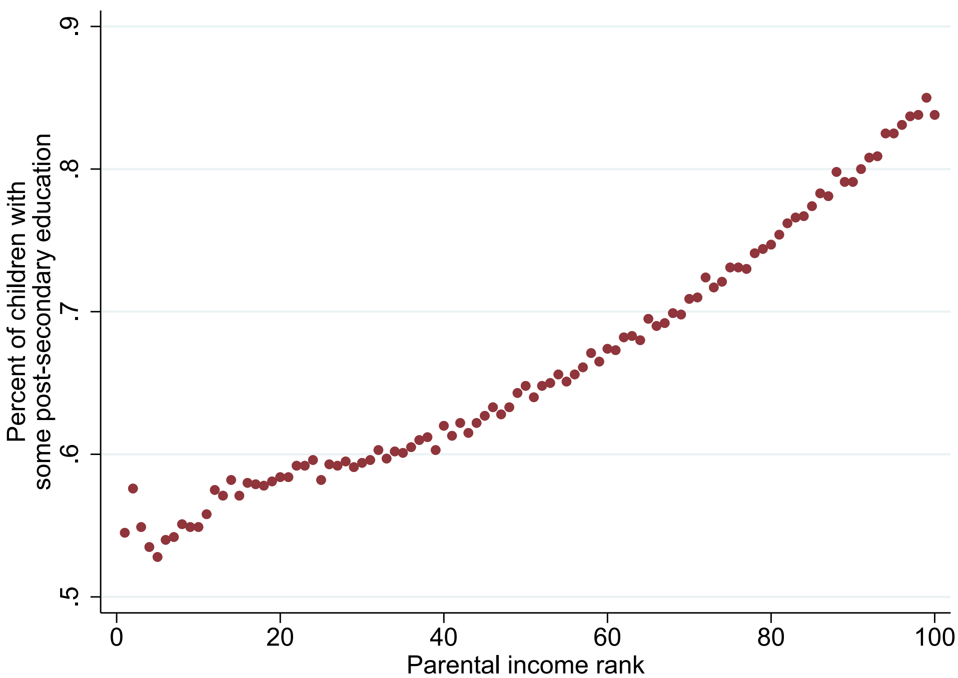

Notes: This figure presents the share of children who ever enrolled in a post-secondary institution in Canada, separately by parental income percentiles.

Figure A7: Average Social Connections Effects by Parental Income Rank

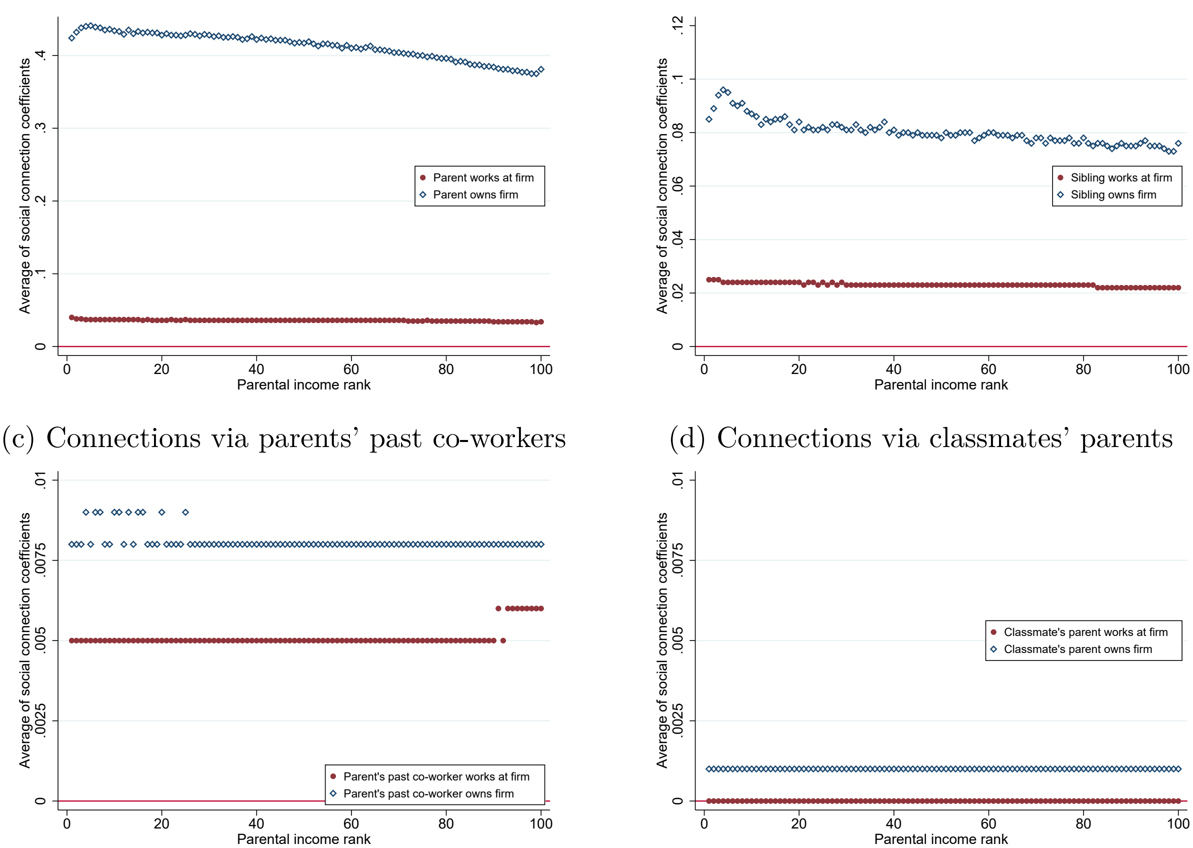

(a) Connections via parents

(b) Connections via siblings

(c) Connections via parents' past co-workers

(d) Connections via classmates' parents

Notes: This figure plots the average social connection effect by parental income rank. That is, we assign each child in our data a value of γ_m^q on the basis of which education category m they belong too, and take the average separately by parental income group. For reasons related to Statistics Canada's vetting procedure, the averages shown on this figure have to be rounded before disclosure. This significantly mutes the amount of variation across parental income groups for connection types shown in panels (c) and (d) since the numbers are very low to begin with.

Figure A8: Distributions of Estimates of Social Connections Effects - by Parent's Job Tenure

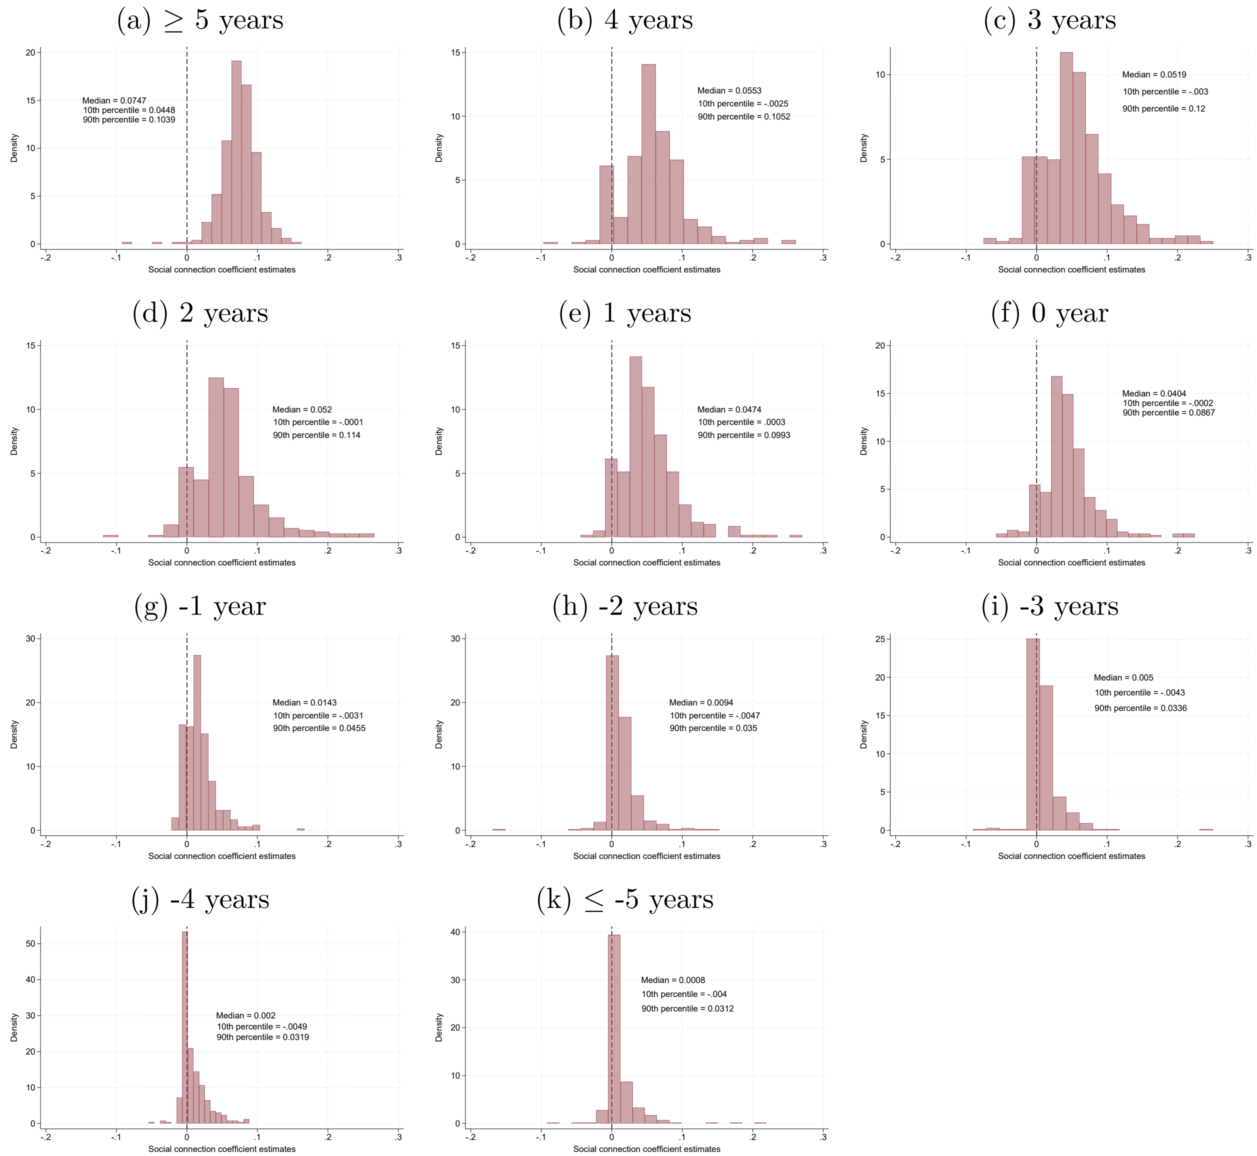

Notes: This figures plots the distributions of social coefficients estimate for different values of job tenure. The estimating equation is $H_{il}^{25} = \alpha_{e(i),l} + \sum_T \gamma_{m(e(i))}^T C_{il}^T + \epsilon_{il}$, where H_{il}^{25} is an indicator for worker i working at firm l when they are aged 25, and C_{il}^T indicates whether a parent has a job tenure of T years at firm l the year their child is aged 25.

Figure A9: Distributions of Estimates of Social Connections Effects - Permutation Tests

(a) Parent works at firm

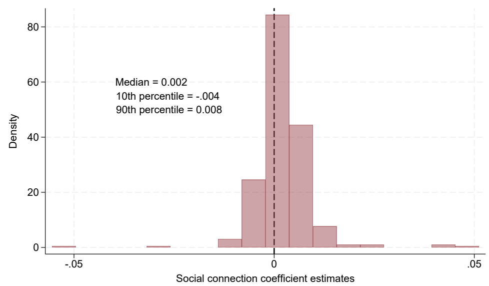

(b) Parent owns firm

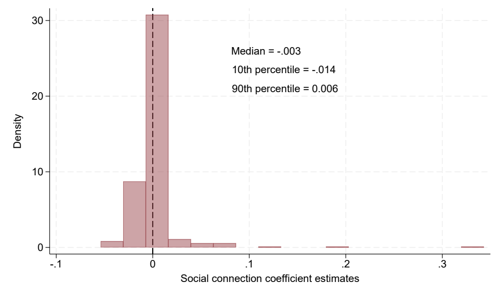

(c) Sibling works at firm

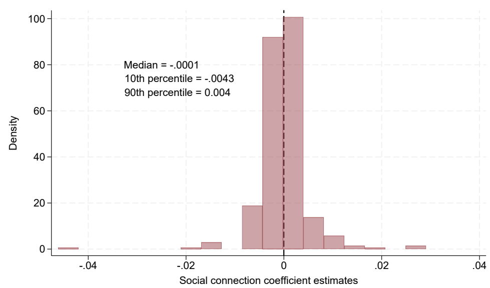

(d) Sibling owns firm

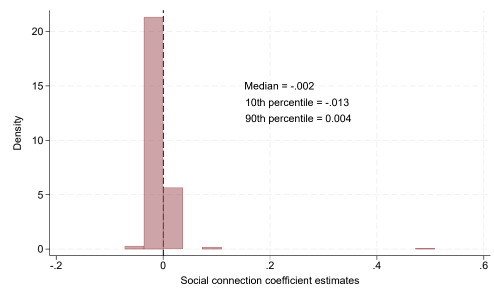

(e) Parent past co-worker works at firm

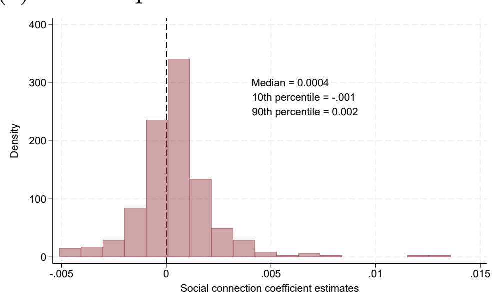

(f) Parent past co-worker owns firm

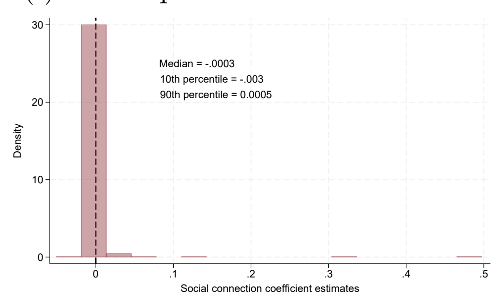

(g) Classmate's parent works at firm

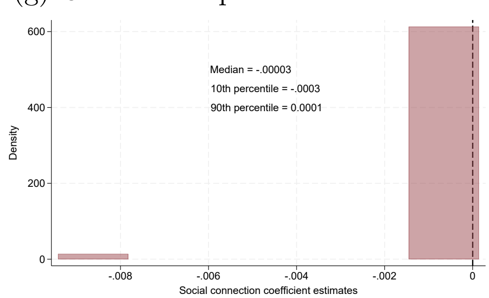

(h) Classmate's parent owns firm

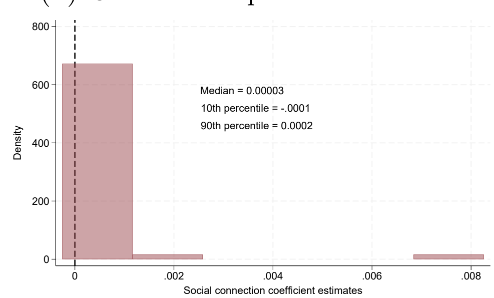

Notes: This figure presents the distribution of estimated placebo social connection coefficients γ_m^q from equation (3). Each panel plots the distribution for a different connection type. In this analysis, the outcome variable $H_{il}^{placebo}$ indicates whether the "control" neighbor of child i works at firm l.

Figure A10: Decomposition of firm sorting component, using unadjusted firm fixed effects

(a) Pooling all connection types

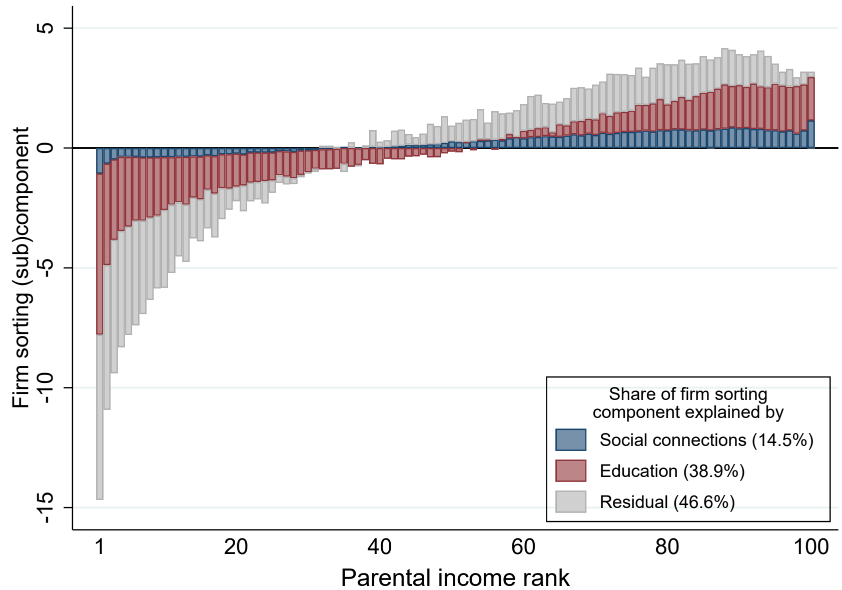

(b) Separately by connection type

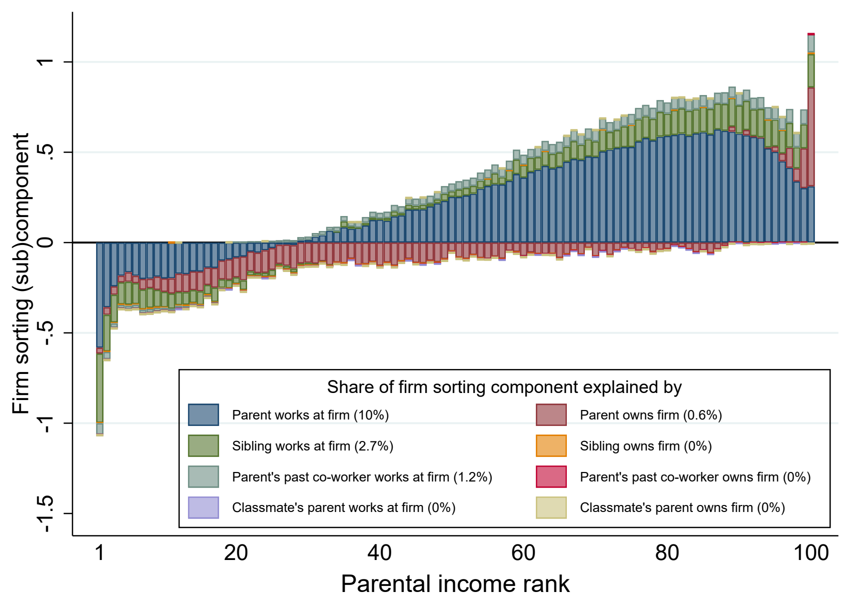

Notes: This figures plots counterfactual firm sorting components based on equations (4) and (5). It uses estimates of δ_l based on equation (8). Panel (a) shows components attributable to education ($\sum_l \delta_l E[\alpha_{e(i),l}|l,p]$) and to social connections pooling all eight types ($\sum_l \delta_l \sum_q E[\gamma^q_{m(e(i))} C^q_{il}|l,p]$), as well as residual unexplained variation in firm sorting. Panel (b) decompose the social connection component, breaking down the contribution of each connection type.

Figure A11: Social Connections and Income Ranks, Singleton Firms Included

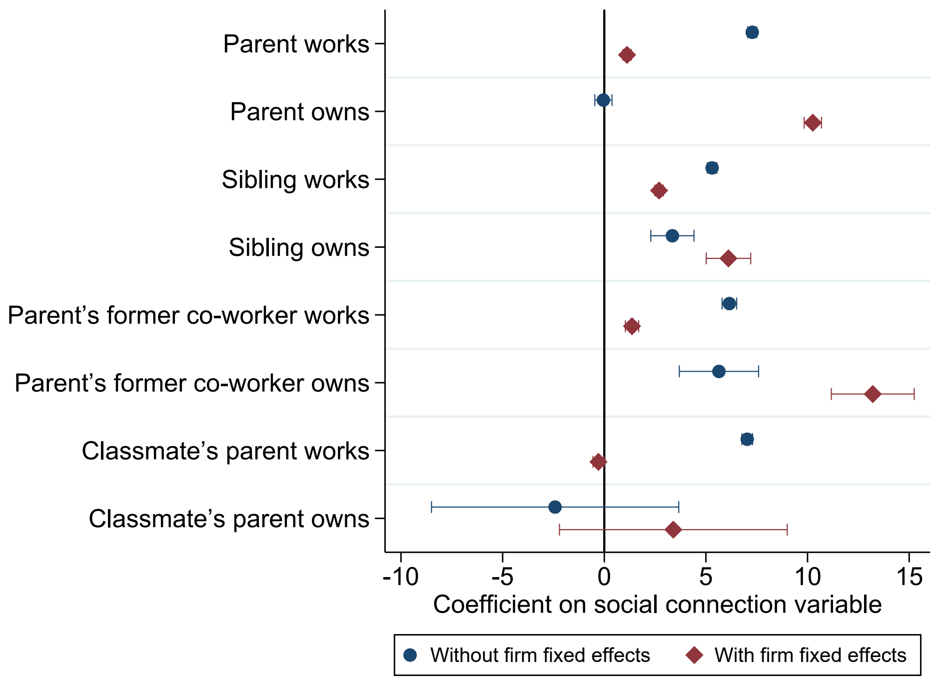

Notes: This figure plots coefficient estimates from equation (6), separately by connection type and by specification that do or do not include firm fixed effects. Singleton firms are pooled into ventiles of the AKM firm effects distribution. Whiskers indicate 95% confidence intervals based on heteroskedasticity-robust standard errors.

Figure A12: Social Connections and Employment Income Ranks

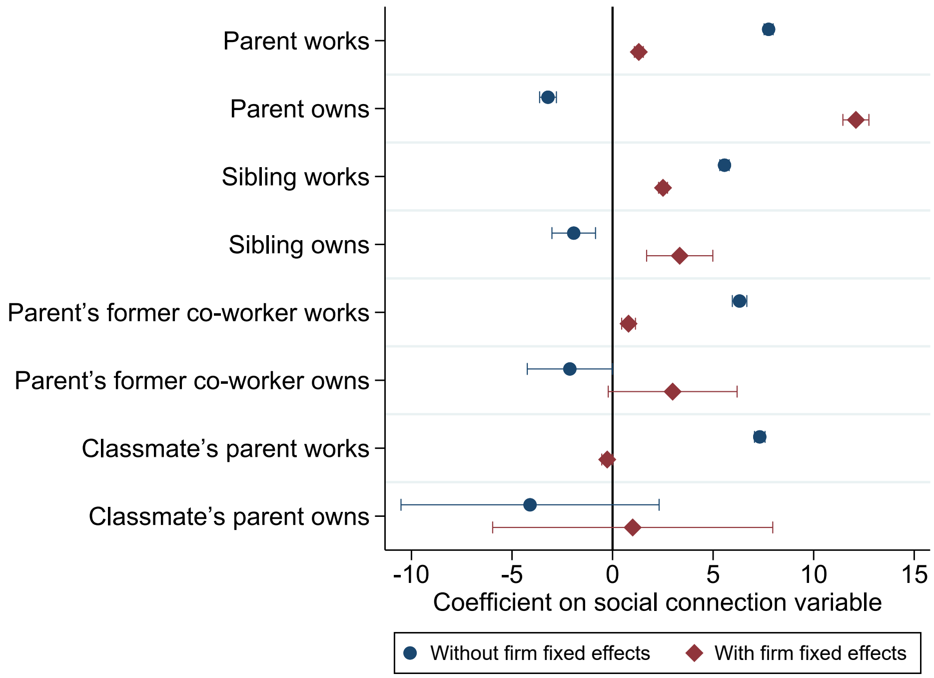

Notes: This figure plots coefficient estimates from equation (6), separately by connection type and by specification that do or do not include firm fixed effects. The outcome is employment income percentile ranks (i.e., excluding non-employment income). Whiskers indicate 95% confidence intervals based on heteroskedasticity-robust standard errors.

Figure A13: Firm Scale and Hiring of Connected Workers

(a) Connections via parents

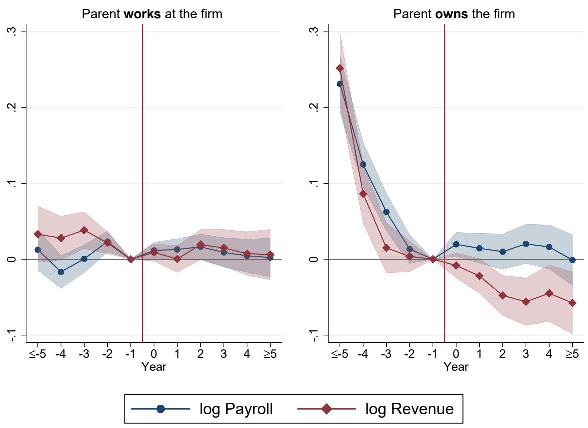

(b) Connections via siblings

(c) Connections via parents' past co-workers

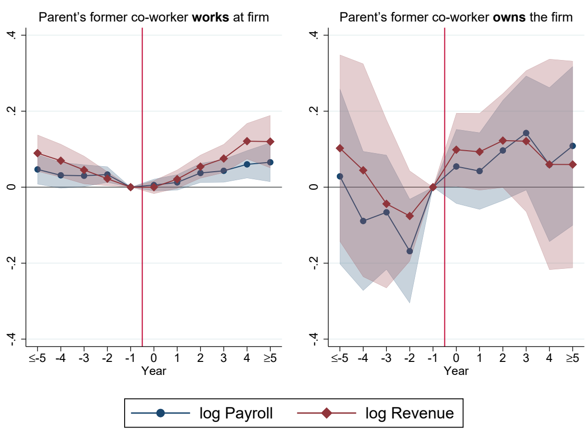

(d) Connections via classmates' parents

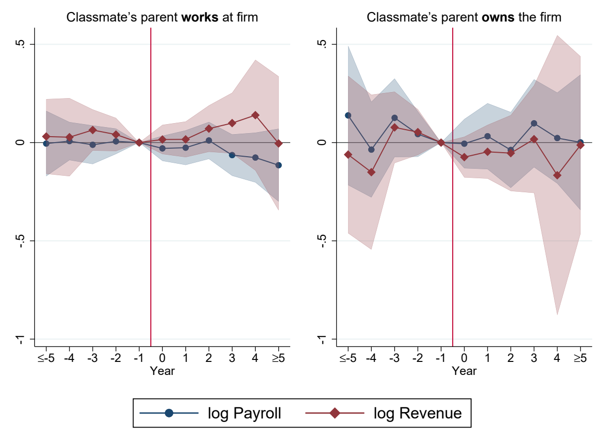

Notes: The figure presents event-study estimates of changes in firm log revenue and log payroll around the time a connected workers is hired. Point estimates are based on equation (7). Shaded areas indicate 95% confidence intervals based on standard errors clustered by comparison set g.

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 (IID) (Corak and Heisz, 1999), which forms the basis for most work on intergenerational mobility in Canada.

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 to the one obtained using the IID (see Figure A4).

A.2 Education Groups

As a first step, we find each person's highest level of education. That is, if an individual holds more than one degree, we select the one at the highest level.[44] Similarly, if an individual enrolled in two different programs but never graduated from either, we assign them to the

[44]For example, if a holder of a Master's degree goes back to school to earn an Associate's degree, we code the Master's degree as their highest level of education even though it is not their most recent program of instruction.

one corresponding to a higher level of education. If the level is the same for both programs, we select the most recent one.[45] At the end of this procedure, each individual who ever enrolled in some post-secondary education is assigned to exactly one program of study.

Having selected each person's most advanced program of study, we then create education groups as unique combinations of a 4-digit CIP code (which correspond to fields of study), a credential type, a program type, a post-secondary institution, and an indicator for having graduated from the program or not. 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 education 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.

[45]For all remaining cases where an individual is assigned two equally recent degrees in the same level of education, we keep the program with the largest enrollment.

B Estimation Notes

B.1 Estimation of Firm Fixed Effects

We obtain estimates of firm fixed effects using the following rank-rank regression

$$y_i = \sum_{p=1}^{100} \beta_p 1\{x_i = p\} + \sum_j \delta_{j(i)} + \varepsilon_i \tag{8}$$

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 β_p coefficients represent the expected income rank of children belonging to parental income group p conditional on the employers they work for, and the estimated δ_j's are firm-level average income ranks adjusted for composition by parental income. Note that this constitutes a regression-compatible decomposition (Fortin, 2008; Haeck and Laliberté, 2025) as β_p is exactly equal to the within-firm component as per equation (1).

Estimation of firm fixed effects in equation (8) is complicated by the fact that there are singleton firms – firms at which we only observe one worker in our sample of children. These firms typically have more than one worker, but only one of them belongs to the three birth cohorts that constitute our estimation sample. Since social connections to 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. To distinguish between high- and low-pay singleton employers, we define the groups using ventiles of the distribution of AKM firm effects.

Figure B1 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, and therefore leaves the rank-rank slope roughly unaffected by the exclusion of singleton firms.

Figure B1: Income rank-rank relationship, with and without children at singleton firms

(a) Unconditional rank-rank relationship

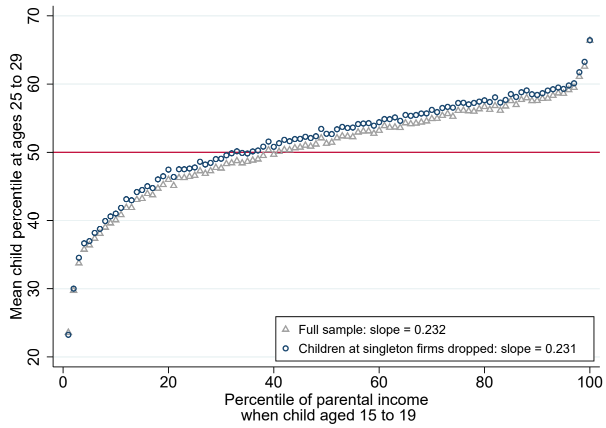

(b) Rank-rank relationship conditional on employers (unadjusted firm effects)

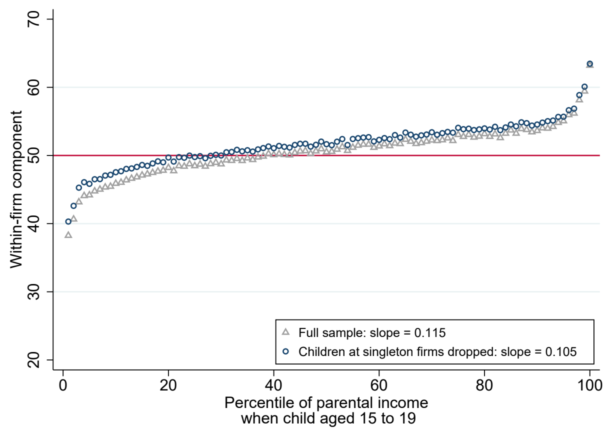

(c) Firm sorting component (unadjusted firm effects)

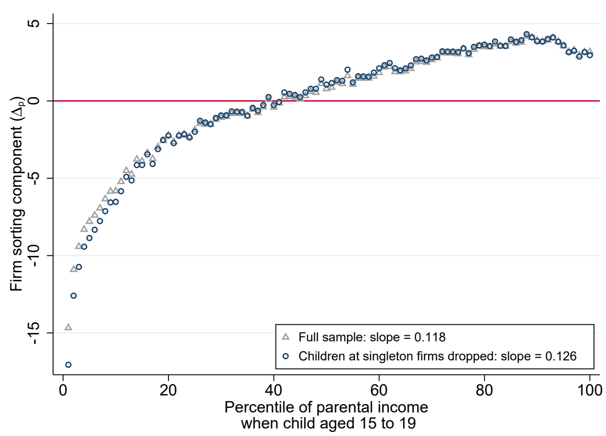

Notes: Panel (a) shows mean child income percentiles for each parental income decile. Grey triangles show unconditional means for the full sample of children born in 1987-89, whereas blue circles show unconditional means for a subsample that excludes children at singleton firms. Panel (b) shows conditional means accounting for differences in employers, as per equation (8). Panel (c) shows the corresponding firm sorting component.

B.2 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 additional sample restrictions. First, we only keep worker-year observations with annual earnings above 5,000$ (in 2016 Canadian dollars). 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. 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 most workers in the sample spent at least one full year in the labor market. Finally, we focus on the largest connected set (Card et al., 2018).

The dependent variable is the log of T4 earnings ($\ln e_{it}$) 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

$$\ln e_{it} = \alpha'_i + \psi'_{j(i,t)} + \beta X_{it} + \varepsilon'_{it} \tag{9}$$

where X_{it} is a vector of covariates, which generally includes 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

$$r_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it} \tag{10}$$

and store the estimated $\hat{\alpha}_i$'s and $\hat{\psi}_j$'s. For validation purposes, we estimated eq. (9) on smaller sample for which estimation is computationally feasible (all 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.3 Selection-adjusted Income-rank Firm Premiums

Here, we cast the problem of obtaining selection-adjusted firm effects in eq. (8) 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. (8) with a cubic polynomial in AKM worker effects.[46] Hence, the firm fixed effects we estimate this way are adjusted for compositional differences across firms in terms of AKM worker effects. Rather than grouping singleton firms like we do when estimating eq. (8), here we drop them from the estimation sample and impute 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 using fitted values from a model in which we project the estimated selection-corrected δ_j's (obtained controlling for AKM worker effects on the set of non-singleton firms) on their corresponding AKM firm effects ψ_j.[47]

Robustness. For robustness purposes, we report results based on an alternative method for adjusting firm effects for selection. Here, we generate counterfactual income ranks for each child i in our sample by imposing that everyone receives the same AKM firm effect ψ_j. 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 effect for child i's modal employer j between the age of 25 and 29, which is in log units. We can then write

$$\ln I_i = a_i + \psi_{j(i)} \tag{11}$$

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

$$\tilde{I}_i = \exp(a_i + \bar{\psi}^c) \tag{12}$$

[46]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. This means including AKM worker effects in the model can account for selection into non-employment between the ages of 25 and 29.

[47]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.

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.

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. But since AKM firm effects don't exist for the two non-employment categories, as well as for firms outside the largest connected set, we must fill these holes using the estimates from the main approach described above.

Figure B2 shows the firm sorting component Δ_p for the two different ways we selection-adjust firm effects. Using the linear rank-rank slope as a summary measure of the firm sorting component, we find that results are fairly stable across measurement approaches. The slope is equal 0.072 when controlling for AKM worker effects and to 0.067 when directly using AKM firm effects.

We note that the main 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.

Figure B2: Robustness to alternative estimates of selection-adjusted firm effects

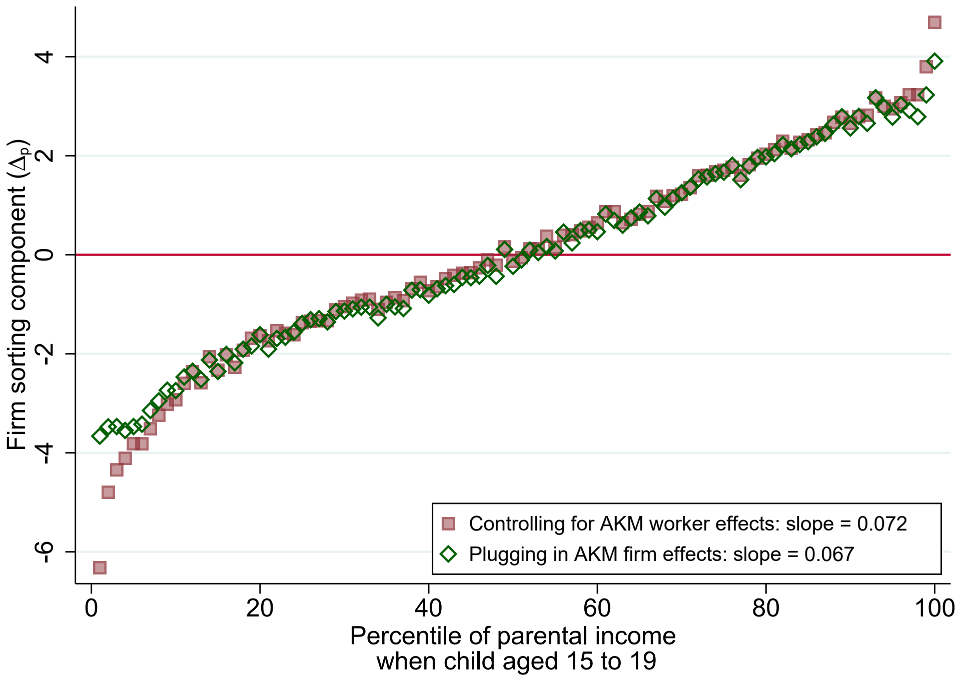

Notes: The series in red squares and green diamonds show alternative sorting components in which selection-adjusted firm fixed effects are plugged in equation (1). The two methods for obtaining selection-adjusted firm fixed effects are described in the text.

B.4 Retrieving Education-firm Fixed Effects

For each category m, the estimating equation for the full dyadic sample (i.e., including cells that do not contribute to identification of the social connection coefficients) is

$$H_{il} = \kappa_m + \alpha'_{e(i),l} + \sum_q \gamma_m^q C_{il} + \epsilon_{il} \tag{13}$$

where we normalize estimates of $\alpha'_{e(i),l}$ such that their within-sample average is zero, and include an intercept κ_m so that $\alpha_{e,l} = \kappa_{m(e)} + \alpha'_{e,l}$. Since estimation of eq. (13) on the full sample is computationally infeasible, we must back-out the key parameters using estimates of γ_m^q as well as other moments of the data.

First, taking the sample average of eq. (13) we obtain

$$\bar{H}_m = \hat{\kappa}_m + \underbrace{\bar{\alpha}'_m}_{=0} + \sum_q \hat{\gamma}_m^q \bar{C}_m^q + \underbrace{\bar{\epsilon}_m}_{=0} \tag{14}$$

where $\bar{C}_m^q$ is the group-m sample average of C_{il}^q. Since every child is matched with one firm and the sample includes all possible worker-firm pairs, $\bar{H}_m = \frac{1}{J}$, where J is the number of firms in the data. This implies we can back-out the value of $\hat{\kappa}_m$:

$$\hat{\kappa}_m = \frac{1}{J} - \sum_q \hat{\gamma}_m^q \bar{C}_m^q. \tag{15}$$

Then, averaging equation (13) across education-firm cells $\{e, l\}$ we obtain

$$\bar{H}_{e,l} = \hat{\alpha}'_{e,l} + \hat{\kappa}_{m(e)} + \sum_q \hat{\gamma}_{m(e)}^q \bar{C}_{e,l}^q. \tag{16}$$

Note that $\bar{H}_{e,l}$ is equal to $s_{l|e}$ – the fraction of workers with education e that actually works at firm l – which can be directly measured. Combining equations (15) and (16),

$$\hat{\alpha}'_{e,l} = \left(s_{l|e} - \frac{1}{J}\right) - \sum_q \left[\hat{\gamma}_{m(e)}^q \left(\bar{C}_{e,l}^q - \bar{C}_{m(e)}^q\right)\right] \tag{17}$$

or, put differently,

$$\hat{\alpha}_{e,l} = \hat{\kappa}_{m(e)} + \hat{\alpha}_{e,l} = s_{l|e} - \sum_q \hat{\gamma}_{m(e)}^q \bar{C}_{e,l}^q. \tag{18}$$

Note that if $s_{l|e} = 0$ (i.e., no one from education group e works at firm l) and $\bar{C}^q_{e,l} = 0$ for all connections types, then $\hat{\alpha}'_{e,l} = -\hat{\kappa}_m$. With estimates of $\alpha'_{e,l}$, κ_m, and γ^q_m in hand, we can then compute predicted values $\hat{H}_{il}$ for the entire dyadic sample

$$\hat{H}_{il} = \left(\hat{\kappa}_{m(e(i))} + \hat{\alpha}'_{e(i),l}\right) + \sum_q \hat{\gamma}^q_{m(e(i))} C^q_{il}. \tag{19}$$

Finally, with some abuse of notation, we take expectations conditional on parental income rank and firm to obtain counterfactual distributions of workers to firms:

$$\tilde{s}_{l|p} = E[\hat{H}_{il}|l, p]. \tag{20}$$

In practice, we calculate $\tilde{s}_{l|p}$ by collapsing over education groups

$$\tilde{s}_{l|p} = \sum_e E[\hat{H}_{il}|l, p, e]\nu_{e|p} \tag{21}$$

$$= \sum_e \hat{\alpha}_{e,l}\nu_{e|p} + \sum_e \left(\sum_q \hat{\gamma}^q_{m(e)} \bar{C}^q_{ep,l}\right) \nu_{e|p} \tag{22}$$

where $\bar{C}^q_{ep,l} = E[C^q_{il}|l, p, e]$ is the average of C^q_{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.

C Sensitivity and Robustness of Rank-rank Decompositions

C.1 Alternative Income Measures

Income measured at older ages. To examine the sensitivity of the rank-rank decomposition to the age at which income and firms are measured, 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 with sufficient accuracy. For benchmarking purposes, in Figure C1 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 C2 shows both the unconditional rank-rank relationship as well as the estimated ranks conditional on main employer (based on eq. (8)) 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 prospects. 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 (based on eq. (8)). The resulting average ranks conditional on employers, shown in Figure C3, 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 main measure of income includes all sources of income. This is important since one's main form of remuneration is itself an outcome of labor market opportunities. For instance, Gendron-Carrier (2023) shows that unincorportated business ownership is a close substitute for employment in low-pay firms. In addition, some firms may directly contribute to workers' non-labor income – e.g. by offering stock options –

Figure C1: Comparison of rank-rank relationships across birth cohorts

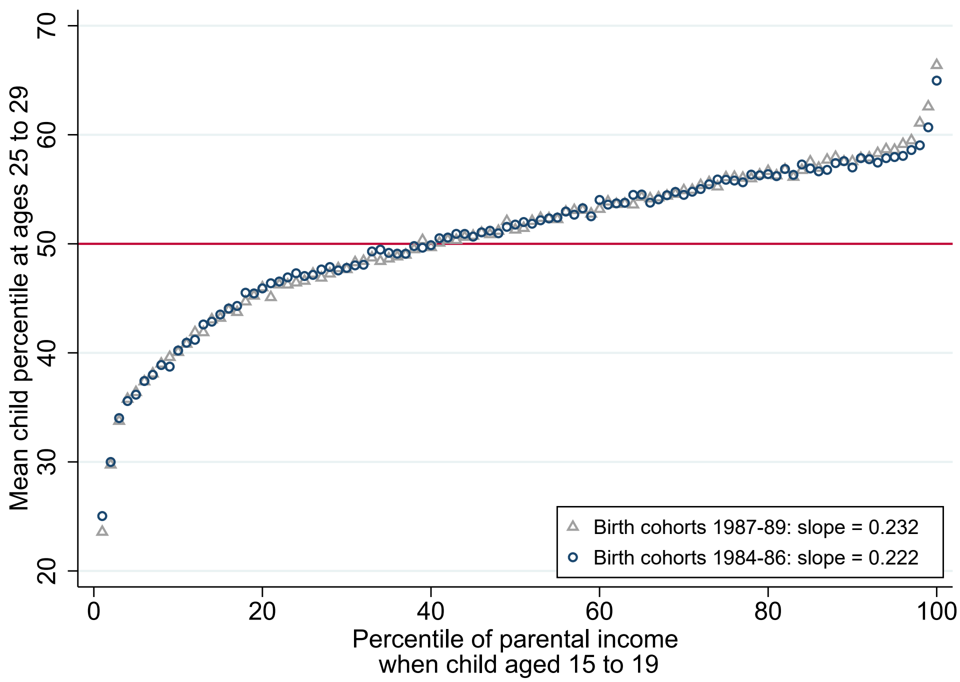

Notes: This figure shows mean child percentile income rank at age 25-29, separately for the 1987-89 and 1984-86 birth cohorts.

Figure C2: Income rank-rank relationship, by age-at-measurement

(a) Average income and main employer at 25-29

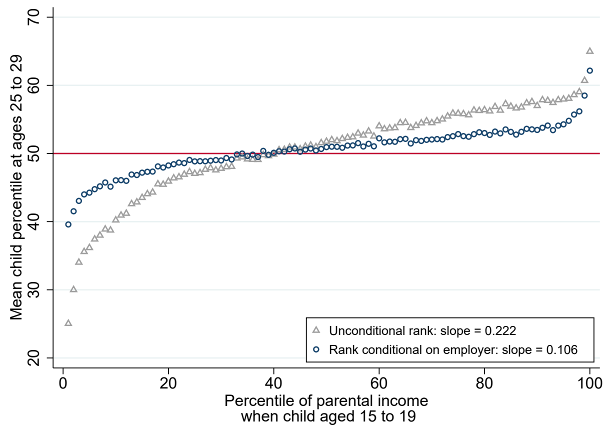

(b) Average income and main employer at 27-31

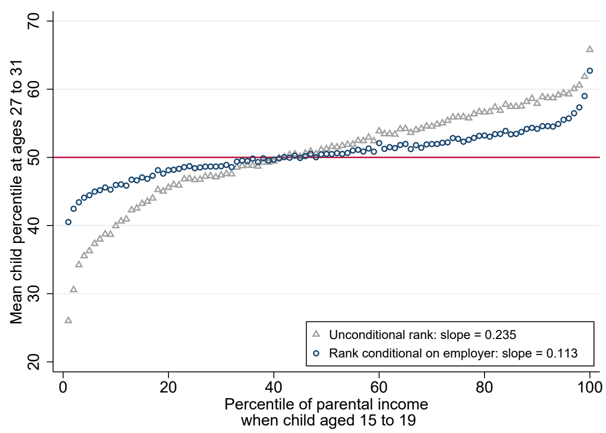

(c) Average income and main employer at 29-33

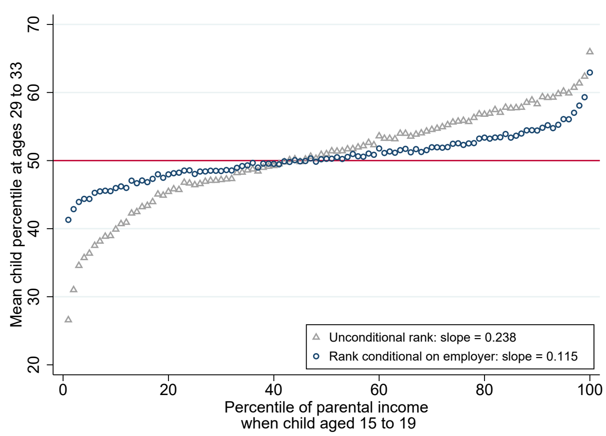

Notes: This figure shows mean child percentile income rank for the 1984-86 birth cohorts. Average ranks conditional on employers are based on equation (8). In panel A the outcome is average income rank based on average income between the ages of 25 and 29. Panels B and C reproduce similar analyses using average income at 27-31 and 29-33, respectively.

Figure C3: Income mobility and past employers

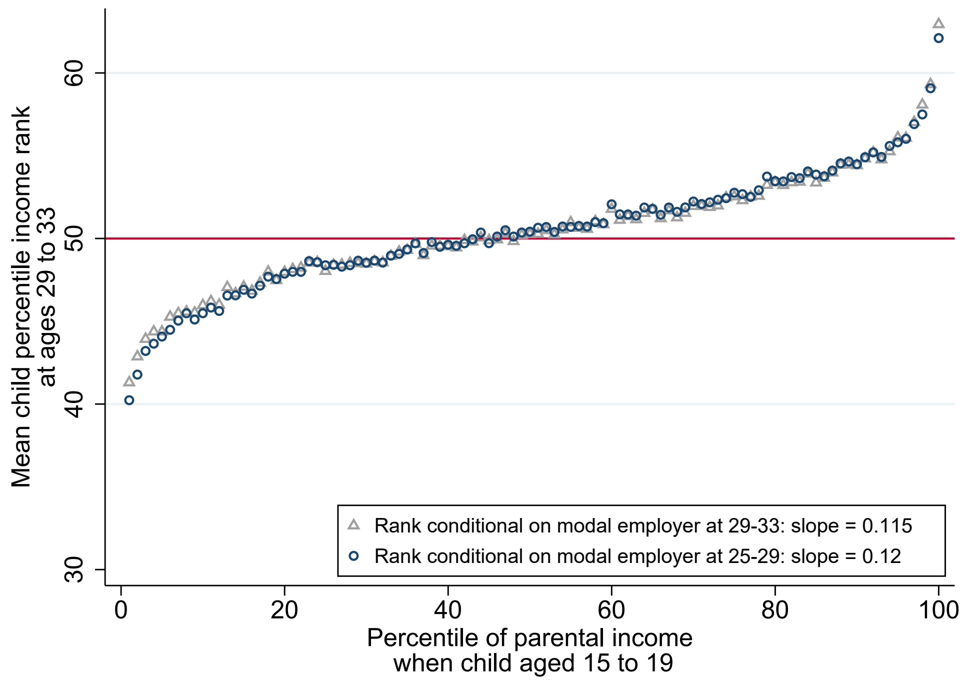

Notes: This figure shows mean child percentile income rank at age 29-33 for the 1984-86 birth cohorts. Both series are conditional on one's main employer, as per equation (8). The series in grey triangle conditions on main employer at 29-33, and the series in blue circles conditions on main employer at 25-29.

and others systematically have high rates of business ownership among their employees. For example, incorporation of one's medical practice is a common phenomenon among physicians in Canada (Strobel et al., 2025). Still, for most people the bulk of the compensation associated with an employment relationship is in the form of wages and salaries. Hence, employers may play a larger role for mobility in terms of labor income than in terms of total income. To examine this, we reproduce our main decomposition analysis using employment income (i.e., T4 earnings) only, in Figure C4.[48]

We find that the unconditional rank-rank slope based on children's employment income alone is 0.204, roughly 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.

[48]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 receives the same amount of non-labor income and so employment income is the only source of variation in income across children.

Figure C4: Intergenerational Employment Income Mobility and Children's Employers

(a) Child-Parent Rank-rank Relationship

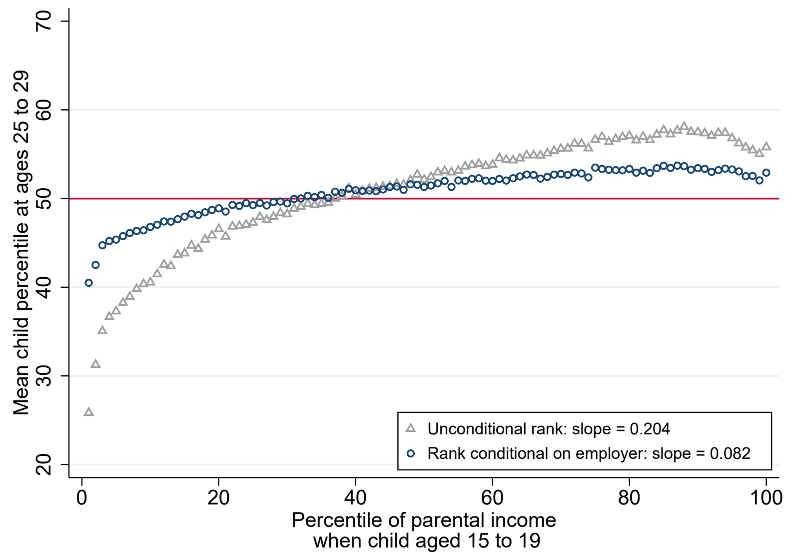

(b) Firm Sorting Component

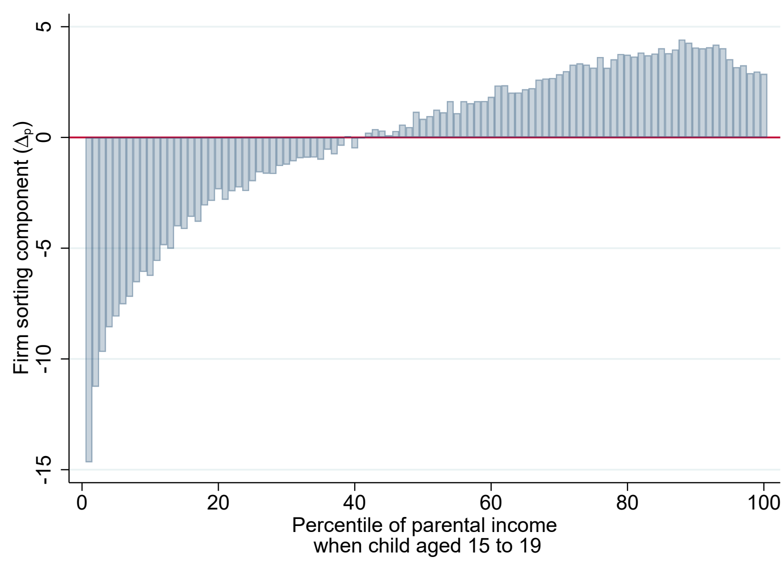

Notes: Panel (a) shows mean child employment income percentiles for each parental income decile. Grey triangles show unconditional means, whereas blue circles indicate conditional means accounting for differences in employers, as per equation (8). The difference in linear slopes represents the contribution of employers to employment income mobility. Panel (b) shows the corresponding firm sorting component. This component is equal to the vertical distance between the two series presented in panel (b).

C.2 Sorting by comparative advantage

While differences in distributions $s_{j|p}$ across parental income groups chiefly influence the sorting component, they may also contribute to within-firm gaps if workers sort by comparative advantage. To examine whether that is the case, we decompose the within-firm component (based on eq. (8)) of the rank-rank income relationship

$$\beta_p = \sum_j \left(\bar{y}_{jp} - \delta_j\right) s_j + \sum_j \left(\bar{y}_{jp} - \delta_j\right) \left(s_{j|p} - s_j\right) \tag{23}$$

where the first term imposes that all groups be equally represented at all firms and therefore is free of comparative-advantage sorting. The second term, then, can be interpreted as a measure of comparative-advantage sorting.

Unfortunately, neither term can be directly calculated due to data sparsity: not all parental-income groups are observed at all firms, and therefore $\bar{y}_{jp}$ is undefined in many cases. As a result, the shares s_j may not add up to one *within the summations*. 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 $\sum_j \left(\bar{y}_{jp} - \delta_j\right) s_j$ won't sum up to one.

Still, the correlation between $\left(\bar{y}_{jp} - \delta_j\right)$ and $\left(s_{j|p} - s_j\right)$ can be informative of the extent of sorting by comparative advantage within the set of observed matches. 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). In our data, that correlation is equal to -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.

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

$$\tilde{\beta}_p = \frac{\sum_j \left(\bar{y}_{jp} - \delta_j\right) s_j}{\sum_j 1\{s_{j|p} > 0\} s_j}. \tag{24}$$

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 C5, we plot the re-scaled sum from eq. (24) against the true β_p. The two series closely overlap throughout the parental income distribution, expect for the top 3 percentiles, where the re-scaled sum slightly mutes the non-linearity. We take these findings as evidence that sorting by comparative-advantage is limited.

Figure C5: Sorting on Comparative Advantage

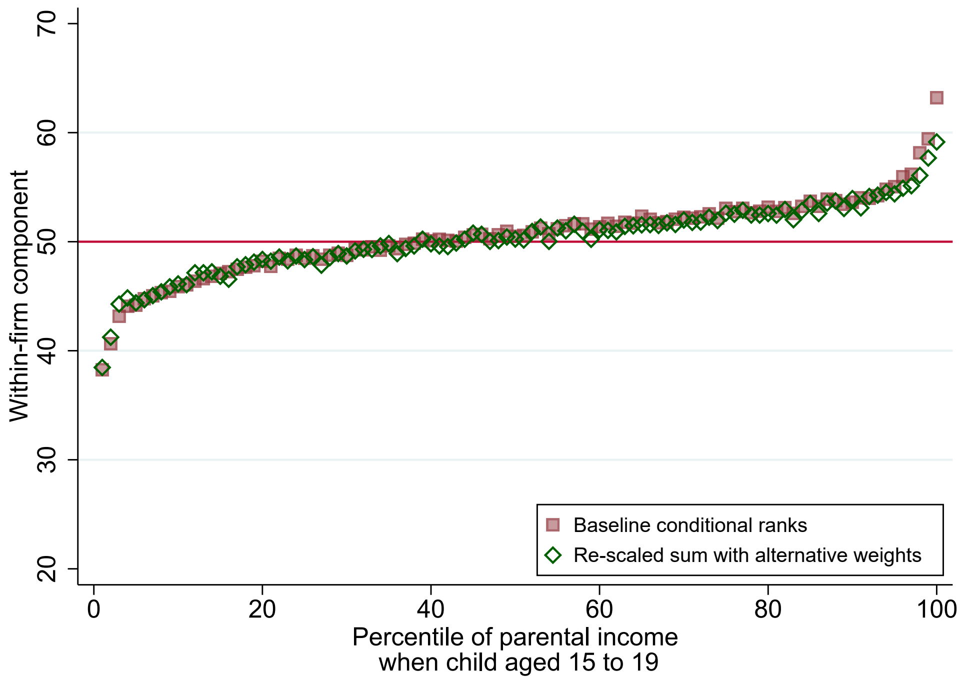

Notes: This figure plots both the within-firm component based on equation (8), and the re-scaled version based on equation (24).

D Decomposing education-firm fixed effects

Patterns of sorting into firms by education may reflect different firms' propensity to hire students from a particular institution or from a particular field of study. They may also capture variation in social connections built in school, either because people who studied together may provide each other with job referrals, or because a professor in a specific program systematically recommends their students to particular firms.

Here, we examine the separate contribution of these channels to the education-based firm sorting component. Since institutions and fields of study are only observed for children with some post-secondary education, we first write

$$\alpha_{e,l} = \alpha_{e,l} PS_e + \alpha_{e,l}(1 - PS_e) \tag{25}$$

where PS_e is an indicator function that takes a value of one for groups e that correspond to some post-secondary level of education. Then, let $\bar{\alpha}_{PS,l} = \frac{\sum_{e \in PS} \alpha_{e,l} \nu_e}{\sum_{e \in PS} \nu_e}$ be the average match effect between firm l and post-secondary education groups of any kind, and let $\bar{\alpha}_{HS,l}$ be the corresponding term for high school education.[49] We can further break down education-firm match effects

$$\begin{aligned} \alpha_{e,l} = & \left[PS_e \left(\bar{\alpha}_{PS,l} - \bar{\alpha}_{HS,l}\right) + \bar{\alpha}_{HS,l}\right] \\ & + \left(\alpha_{e,l} - \bar{\alpha}_{HS,l}\right)\left(1 - PS_e\right) \\ & + \left(\alpha_{e,l} - \bar{\alpha}_{PS,l}\right) PS_e \end{aligned} \tag{26}$$

where the first line captures firm l's propensity to hire workers with post-secondary education relative to workers with high school education. The next two lines are deviations from broad averages. The second-term, $\left(\alpha_{e,l} - \bar{\alpha}_{HS,l}\right)(1 - PS_e)$, captures differences in conditional match probabilities among people who do not have post-secondary education but attended a different set of high schools. The third term,$\left(\alpha_{e,l} - \bar{\alpha}_{PS,l}\right) PS_e$, reflects the strength of the match effect between education group e and firm l relative to that of the average post-secondary program. We further decompose that within-post-secondary education deviation term using the following regression:

$$\alpha_{e,l} - \alpha_{PS,l} = \eta_{institution(e),l} + \zeta_{field(e),l} + \mu_{e,l} \tag{27}$$

[49] If social connections were not included in eq. (2), then $\bar{\alpha}_{PS,l}$ would simply be the conditional probability of matching with firm l if one has some post-secondary education of any kind.

where $\eta_{institution(e),l}$ are post-secondary institution-by-firm fixed effects, $\zeta_{field(e),l}$ are field-of-study-by-firm fixed effects, and fields of study correspond to 4-digit CIP codes. These two components respectively capture the overall propensity of workers from a given institution to work at firm l – e.g., due to geographic proximity between a school and a firm, or because of reputation effects – and the propensity of workers with education in a given field to work at that firm. For example, $\zeta_{field(e),l}$ accounts for the fact that engineering firms are more likely to hire workers with a degree in engineering than from other fields. Finally, $\mu_{e,l}$ is a residual match effect, net of both geographic and broad academic factors. That residual term may capture a potential social connection effect, but may also reflect the fact that an education program could tailor its curriculum for a specific firm's needs.

Finally, we evaluate the economic significance of each of these channels for intergenerational mobility using the following decomposition:

$$\begin{aligned}\tilde{\Delta}_p^{educ} = \sum_l \delta_l \sum_e \alpha_{e,l}\nu_{e|p} &= \sum_l \delta_l \sum_e \left[PS_e\left(\bar{\alpha}_{PS,l} - \bar{\alpha}_{HS,l}\right) + \bar{\alpha}_{HS,l}\right]\nu_{e|p} \\ &+ \sum_l \delta_l \sum_e \left[\left(\alpha_{e,l} - \bar{\alpha}_{HS,l}\right)\left(1 - PS_e\right)\right]\nu_{e|p} \\ &+ \sum_l \delta_l \sum_e \left[\left(\eta_{institution(e),l} + \zeta_{field(e),l} + \mu_{e,l}\right) PS_e\right]\nu_{e|p} \qquad (28)\end{aligned}$$

Results are presented in Figures D1 (using selection-adjusted firm effects) and D2 (using unadjusted firm effects). In panel (a), we show the three broad education-based sorting component, and report the share of the total firm sorting (Δ) they account for. In Figure D1, shares sum up to 45% (up to rounding), which is the share of the firm sorting component attributable to education (see Figure 5). Differences in representation in post-secondary education of any kind across parental income groups accounts for 18% of the firm sorting component. Differences in post-secondary programs conditional on pursue some higher education has the most important explanatory power, accounting for 20% of the firm sorting component. The explanatory power of differences across education groups conditional on not enrolling in higher education is 8% in total, but is particularly important for the bottom 2 parental income percentiles.

In panel (b), we break down the sub-component attributable to differences within the broad set of children with post-secondary education. The most important dimension is the institutions children attend, which partly captures "vertical" differences between colleges and universities. Residual match effects – which are mostly likely to capture social connections – explain 5% of the firm sorting component, which is greater than the explanatory power of fields of study (4%).

Figure D1: Decomposition of firm sorting component by education sub-components, Selection-adjusted firm effects

(a) Broad sub-components

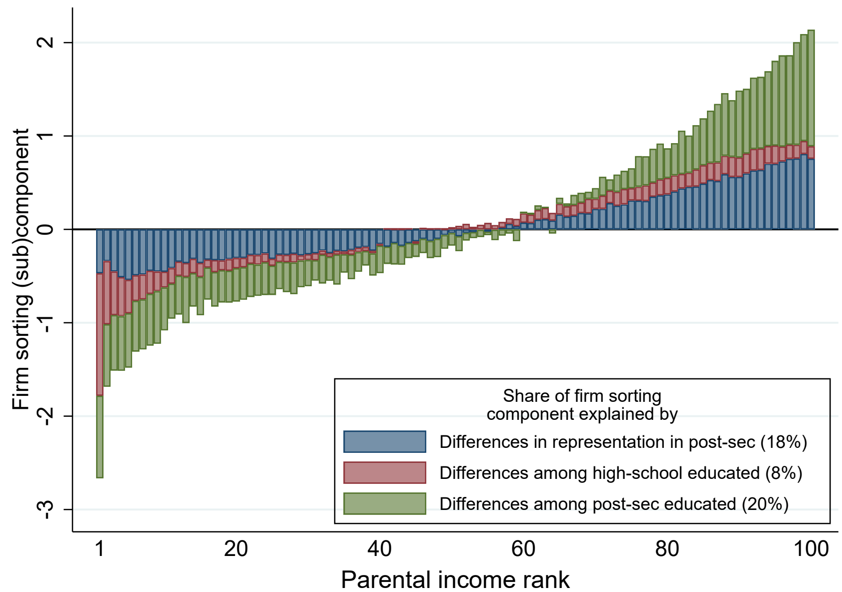

(b) Unpacking differences across post-secondary programs

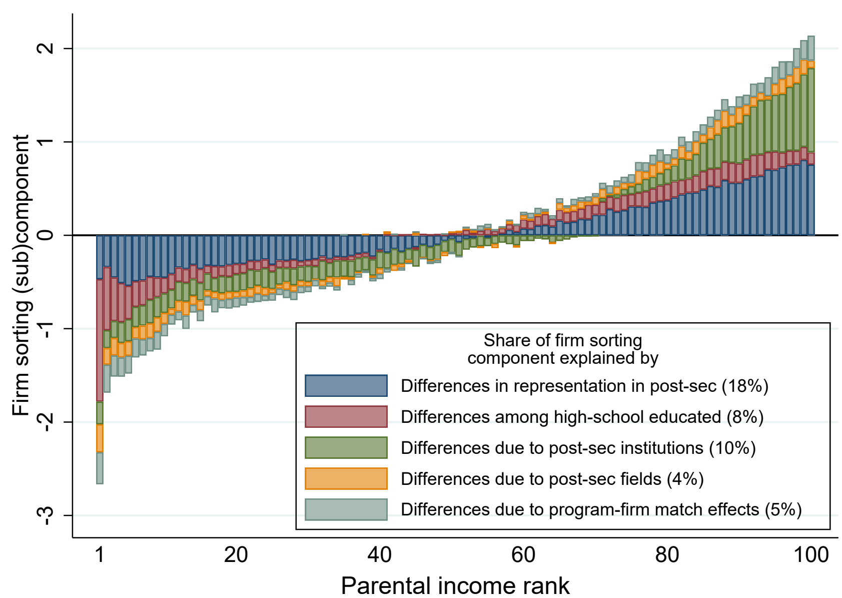

Notes: The figure presents a decomposition of the firm sorting component into sub-components of education. It plots each term in equation (28), using selection-adjusted firm effects as measures of δ_l.

Figure D2: Decomposition of firm sorting component by education sub-components, Unadjusted firm effects

(a) Broad sub-components

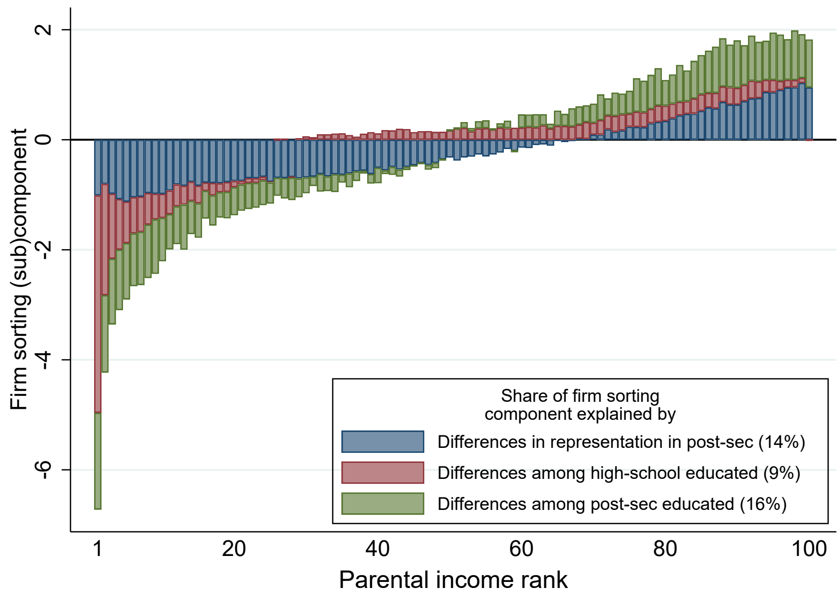

(b) Unpacking differences across post-secondary programs

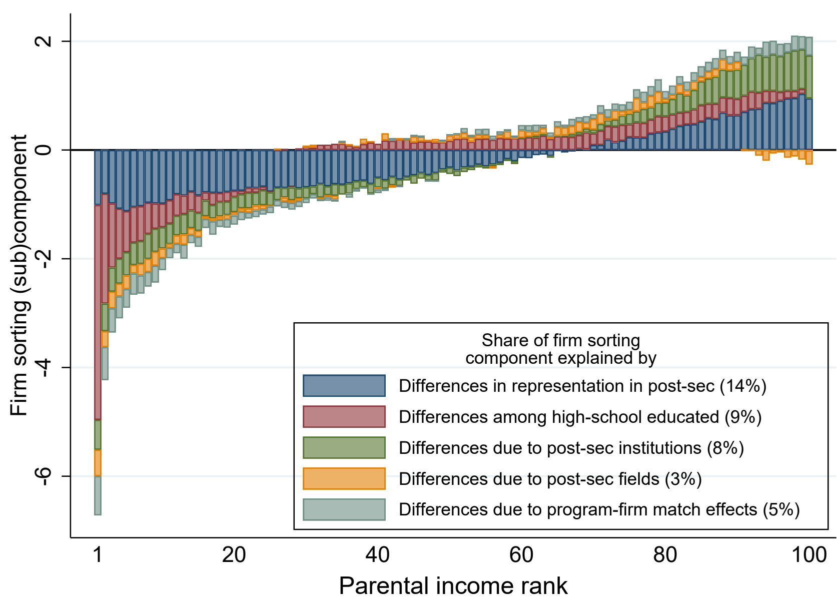

Notes: The figure presents a decomposition of the firm sorting component into sub-components of education. It plots each term in equation (28), using unadjusted firm effects based on equation (8) as measures of δ_l.