Returns to Skills, Technological Change, and Intergenerational Mobility

Jinpeng Shi*

November 5, 2026

Abstract

Does technological change exacerbate or mitigate the intergenerational transmission of economic advantage? Motivated by Tinbergen's race between education and technology, this paper examines how occupational skill requirements, their returns, and intergenerational sorting jointly shape income persistence. We combine linked generations in the Panel Study of Income Dynamics with occupational descriptors from the Dictionary of Occupational Titles and O*NET to construct cognitive, noncognitive, and manual skill requirement indices. Holding occupational requirements fixed, we decompose the intergenerational elasticity of income into skill returns and parental-income gradients in occupational requirements. Cognitive requirements account for the largest positive measured skill contribution. Their rising income returns are partly offset by intergenerational sorting, leaving cognitive contributions and aggregate intergenerational mobility unchanged across the 1950s and 1980s birth cohorts. We then introduce historical skill requirement profiles and further partition income persistence into baseline persistence, skill reevaluation, within-occupation skill requirement evolution, and between-occupational sorting. The measured within-occupation contribution remains small across detailed and broader occupational classifications. Occupational cognitive requirement rankings remain relatively stable despite measured increases in requirement levels over time. Together, these findings explain why rising cognitive demands and skill premia coexist with relatively stable intergenerational income persistence.

*Department of Economics, George Mason University.

Acknowledgments: The author received no financial support for this research and declares no financial relationships or other potential conflicts of interest related to this work.

1 Introduction

Changes in the rewards to skills reshape income inequality, but their implications for the persistence of family advantage depend on who reaches the occupations in which those skills are rewarded. A larger income advantage associated with cognitive-intensive work can strengthen intergenerational persistence when access is concentrated among children from affluent families. The same change can coexist with relatively similar persistence if occupational access becomes less strongly associated with parental resources. Occupational requirements may also evolve without a corresponding change in either occupational rankings or the association between family background and adult income. Distinguishing these margins is essential for connecting changes in the skill content of work to intergenerational mobility. This question connects evidence on skill demand and occupational change (Katz and Murphy, 1992; Autor et al., 2003; Autor and Dorn, 2013) to research on the transmission of economic advantage (Becker and Tomes, 1986; Nybom and Stuhler, 2024).

The empirical distinction matters because occupational requirements, individual skill formation, and family-background differences in access are related but separate margins. Hermo et al. (2022) connect changes in relative labor-market returns to the evolution of different cognitive skills, while Cunha et al. (2010) model the development of cognitive and noncognitive skills through parental environments and investments. Spitz-Oener (2006) documents changes in requirements within occupations. These findings motivate studying both changing rewards and changing job content. Our objective is to measure how these occupational margins enter intergenerational income accounting.

We extend the logic of Tinbergen's race between education and technology to intergenerational allocation (Tinbergen, 1974; Goldin and Katz, 2008). Technological change can increase skill demand and skill premia, potentially strengthening persistence when access to rewarded skills remains concentrated among advantaged families. Occupational reallocation can offset this force if children from disadvantaged backgrounds make larger gains relative to their parents' occupations. Over time, unequal opportunities for skill investment may reinforce family-background differences in access. This is the mechanism motivating the paper, rather than a predetermined prediction that higher productivity must increase or decrease mobility. Our proposition formalizes the ambiguity, and the empirical decompositions examine its price and allocation margins.

This paper develops an occupational-requirement accounting of intergenerational income persistence in the United States. We combine parent-child links from the Panel Study of Income Dynamics with occupational descriptors from the Dictionary of Occupational Titles and O*NET. The PSID provides a long span of linked income observations, but not a common

battery of individual cognitive and noncognitive assessments across these generations. These descriptors allow us to study the work into which children from different family backgrounds sort. They are imputed characteristics of jobs, however, rather than direct observations of workers' skill endowments. The resulting income coefficients are conditional associations with occupational requirements.

The analysis proceeds in two stages. First, we hold occupational requirements fixed and construct cognitive, broad noncognitive, and manual indices. We decompose the intergenerational elasticity of income into the income association of each requirement index and its parental-income gradient. Second, we assign historical profiles to both generations and partition skill requirement-related persistence into baseline persistence, skill reevaluation, within-occupation change, and between-occupation change, retaining an unexplained remainder. The historical analysis begins with three snapshots and then uses multiple dated cognitive profiles, allowing requirements to change even when the underlying Census occupational coding regime is unchanged. A further exercise constructs profiles for 22 occupational families to examine sensitivity to aggregation.

Annual adult child income is observed at ages 30-35, while parental log income is averaged over available observations when the child was aged 15-17. The use of income windows also reflects the measurement concerns emphasized by Solon (1992) and Haider and Solon (2006); age controls and short averages do not guarantee that observed income equals lifetime income. The pooled child-year specification follows Nybom and Stuhler (2024). We report separate birth-decade estimates and retain the original father-first, mother-fallback selection rule. The youngest observations belong to the 1990s birth-decade group but cover only ages 30-32; we therefore treat their age-33 coefficients as provisional. Main interpretation emphasizes the cohorts with more complete age support, and all decomposition equations within a specification use the same observations.

Three findings emerge. First, cognition accounts for the largest positive measured skill component in the fixed-profile benchmark. Between the 1950s and 1980s cohorts, its conditional income coefficient rises from 0.212 to 0.351, while its parental-income gradient falls from 0.488 to 0.310. Their products are consequently similar at these endpoints, 0.104 and 0.109, although intermediate cohorts differ. This pattern illustrates how changes in occupational returns and family-background sorting can offset one another without establishing a monotonic trend or a causal skill supply response. Noncognitive and manual requirements have smaller net contributions, but separate subdomains reveal offsetting associations that broad aggregates can conceal.

Second, our dynamic cognitive decomposition does not uncover a large within-occupation contribution to income persistence. Detailed-occupation estimates range from -0.0028 to

0.0030, and the 22-family exercise preserves small magnitudes even though some cohort-specific confidence intervals exclude zero. Direct profile diagnostics show persistent occupational rankings, while we observe a secular cognitive task upgrading.

Third, an offsetting between-occupation contribution coexists with substantial inequality in access to cognitive demanding occupations. In the 1980s cohort, under the child-period skill requirement ranking, 11.8 percent of children from the lowest parental-income quintile reach the highest cognitive occupational group, compared with 45.3 percent from the highest quintile. Raw upward-movement rates also depend strongly on parental starting positions: the decline between the 1950s and 1980s cohorts is much smaller after standardizing origins. Children can make larger gains relative to a disadvantaged origin while remaining less likely to reach a favorable destination. The transition evidence thus gives an interpretation to the income accounting that aggregate persistence alone cannot provide.

The paper's empirical contribution is to connect intergenerational income accounting to multidimensional occupational requirements, their historical measurement, and observed occupational destinations. This approach extends the information available in long family panels while making the source and interpretation of the imputation explicit.

The remainder of the paper is organized as follows. Section 2 relates the analysis to the literature. Section 3 presents the cross-country motivation and the proposition on skill returns and intergenerational sorting. Section 4 connects worker skills and job tasks to income persistence. Section 5 describes the PSID sample, occupational crosswalks, and skill measures, and Section 6 presents the decomposition framework. Section 7 reports the fixed-profile results. Section 8 introduces historical requirements and their decomposition, and Section 9 examines occupational movement and access by family background. Section 10 discusses sensitivity and the scope of the evidence. Section 11 concludes. The appendix provides construction details, additional estimates, and labor-income results.

2 Related Literature

This paper brings together research on technological change and the wage structure with research on the intergenerational transmission of economic advantage. In the race between education and technology, changes in relative skill supply and demand jointly determine skill premia (Tinbergen, 1974; Katz and Murphy, 1992; Goldin and Katz, 2008). Autor et al. (2003) and Acemoglu and Autor (2011) connect this framework to the changing task content of work. Spitz-Oener (2006) documents changes in requirements within occupations, while Autor and Dorn (2013) examine occupational reallocation associated with computerization. We bring an intergenerational allocation margin to this framework: the implications of changing

skill rewards depend on which children enter the occupations receiving those rewards. Our analysis distinguishes changes in occupational requirements from changes in their income returns and relates both to parental background.

The closest connection is to research explaining why changes in skill premia need not produce corresponding changes in intergenerational income persistence. Nybom and Stuhler (2024) show that changes in the parental generation can partially offset the implications of rising skill premia for US mobility. Their analysis establishes the importance of intergenerational dynamics when interpreting contemporaneous mobility trends. We examine these relationships through occupational requirements and destinations. Holding requirements fixed separates their income returns from their parental-income gradients; introducing historical profiles further distinguishes baseline persistence, skill reevaluation, and within- and between-occupation changes. This makes it possible to examine whether the occupational contribution to persistence reflects changing rewards, changing job content, or movement between parental and child occupations.

Our analysis also relates to the literature on multidimensional human capital and the transmission of family advantage. Becker and Tomes (1986) and Becker et al. (2018) emphasize parental investment, while Cunha et al. (2010) estimate the formation of cognitive and noncognitive skills. Blanden et al. (2007) account for income persistence through skills, education, and labor-market attachment. Research on skill returns further demonstrates the importance of distinguishing dimensions of human capital: Lindqvist and Vestman (2011), Deming (2017), and Edin et al. (2022) examine cognitive, social, and noncognitive capacities, and Hermo et al. (2022) connect changing relative returns to the evolution of different cognitive skills. We complement studies using individual assessments with occupation-imputed requirements that can be linked to families across a long historical span. These measures describe the work individuals enter, allowing us to study occupational allocation across generations.

A related literature studies occupations as a channel of intergenerational persistence. Haeck and Laliberté (2025) document occupational segregation by parental income and distinguish access to favorable occupations from following a parent's occupation. We extend this distinction to the cognitive requirements of occupational origins and destinations, allowing requirement profiles to vary historically. Our decomposition measures the parental-income gradient of income-valued occupational differences, while transition matrices describe destinations conditional on family background and occupational origin. These exercises clarify why an offsetting between-occupation contribution can coexist with unequal access to cognitively demanding work: gains relative to parental origins and access to favorable destinations are different margins.

Finally, the paper contributes to the discussion of technology, inequality, and mobility underlying the Great Gatsby curve. Corak (2013) documents the cross-country relationship between inequality and persistence, and Durlauf et al. (2022) review its possible mechanisms. Hassler et al. (2007) show that different sources of inequality can generate different relationships with mobility. Theoretical work by Galor and Tsiddon (1997) and Hassler and Rodríguez Mora (2000), and evidence on computerization from Arntz et al. (2025), directly connect technological change to intergenerational allocation. Our contribution is to examine the occupational margins through which changing rewards can reinforce or offset family-income persistence. The cross-country TFP and PISA comparisons motivate the timing and family-background dimensions of this question; the linked-family analysis provides a more detailed accounting of occupational requirements, returns, and movement.

3 The Great Gatsby Curve, Technology, and Intergenerational Skill Allocation

3.1 Technological Change and the Intergenerational Race for Skills

The Great Gatsby curve associates greater income inequality with stronger intergenerational persistence. Figure 1 reproduces this relationship using a mobility measure defined as $1 - \text{IGE}$. Combined with evidence that skill-biased technological change raises skill premia, the curve suggests that technological progress could strengthen the transmission of economic advantage. That implication, however, depends on how access to the newly rewarded skills and occupations is distributed across family backgrounds. An increase in the rewards to skills can widen income differences without proportionately strengthening the relationship between parental and child income.

Tinbergen's race provides a useful starting point. When technological change increases skill demand faster than skill supply expands, skill premia rise. In an intergenerational setting, the relevant supply response includes both the development of skills and the allocation of children across occupations. New demand may create opportunities for children to enter occupations different from those of their parents. If children from disadvantaged origins make sufficiently large gains, occupational reallocation can offset the persistence associated with rising skill rewards. Conversely, if advantaged families are better positioned to prepare their children for the newly rewarded work, the response can reinforce family-background differences. The intergenerational implication therefore depends on the distribution and timing of adjustment, rather than on the skill premium alone.

The cross-country comparisons motivate these competing forces. Figure 3 shows a positive

relationship between childhood-period TFP growth and intergenerational mobility. Figure 2 shows that faster TFP growth is also associated with a stronger parental-background gradient in PISA scores, measured net of school effects. Figure 4, which uses adult-period TFP growth, displays a negative relationship with mobility. Together, these patterns motivate a hypothesis in which technological change initially opens occupational opportunities, while subsequent differences in families' capacity to develop rewarded skills can limit the persistence of those gains. The PISA relationship is consistent with a family-background dimension to skill supply, although it does not directly measure parental investment.

The PSID analysis examines the occupational implications of this argument. In the fixed-requirement benchmark, the income return to cognitive requirements is higher for the 1980s cohort than for the 1950s cohort, while the parental-income gradient in those requirements is lower. The two forces leave similar cognitive contributions to IGE at these endpoints. The historical decomposition adds evidence about occupational origins and destinations: an offsetting between-occupation contribution can arise when children from lower-income backgrounds make larger income-valued gains relative to their parents' occupations. This pattern is consistent with occupational reallocation providing a counterweight to the persistence associated with skill rewards.

Such gains do not imply equal access. The transition matrices show that children from affluent families remain substantially more likely to reach the highest cognitive occupational group. Nor does the small within-occupation IGE component imply that job content is unchanged: a change in requirements contributes to persistence only to the extent that its income value varies systematically with parental income. These distinctions motivate the decomposition that follows. The central question is how changing skill prices interact with intergenerational allocation: whether children's occupational movement offsets the advantages attached to parental origins, and how unequal access shapes that adjustment.

3.2 A race between skill demand and skill supply

For the two-skill illustration, let c and n denote cognitive and non-cognitive skills, respectively, and let p and o denote the parent and offspring generations. The income equations in the original formulation are

$$y^p = \tau_c^p c^p + \tau_n^p n^p + \varepsilon^p, \tag{1}$$

$$y^o = \tau_c^o c^o + \tau_n^o n^o + \varepsilon^o. \tag{2}$$

3.3 Exogenous shocks on intergenerational mobility

Assume a mass of homogeneous workers endowed with the same set of skills. Throughout the paper, we remain silent on the accumulation process of skills since skills can be genetically and culturally transmitted within families and acquired through (public) investment [1]. We interpret skill endowment as the stock of skills prior to labor market entry. We only allow 1 generation to stay in the labor market and ignore the timing of reproduction.

Suppose an exogenous shock (e.g., technological change, war, institutional transition, etc.) shifts the (relative) returns to skills across 2 generations, and the offspring generation has the same skill endowment as the parent generation but differs on skills supply to a degree of $1 - \rho_j$. We interpret ρ_j as intergenerational persistence of skill supply. We will introduce endogenous skill supply when considering the heterogeneity of workers in the next section. We further impose a covariance constraint $cov(c, n) = corr(c, n) = \gamma$, on 2 types of general skills within and across generations. Thus, the implication of skills re-evaluation in the labor market on intergenerational elasticity estimates ($\tilde{\beta}^{IGE}$, with shocks) is shown by the next theorem.

Assumption 1. *No serial correlation between error terms across generations,* $\varepsilon^o \perp \varepsilon_p$.

Proposition 1. *Under Assumption 1, the effect of technological progress on the level of intergenerational mobility is ambigous.*

When technological progress alters the relative returns to skills within the labor market, the effect on the level of intergenerational mobility depends on the original magnitude of returns to skills and intergenerational persistence of skill supply.

The original derivation and comparative statics are reproduced in Appendix E.

[1] Modeling the accumulation process of skills is beyond the scope of the paper, but it is an interesting area for future research.

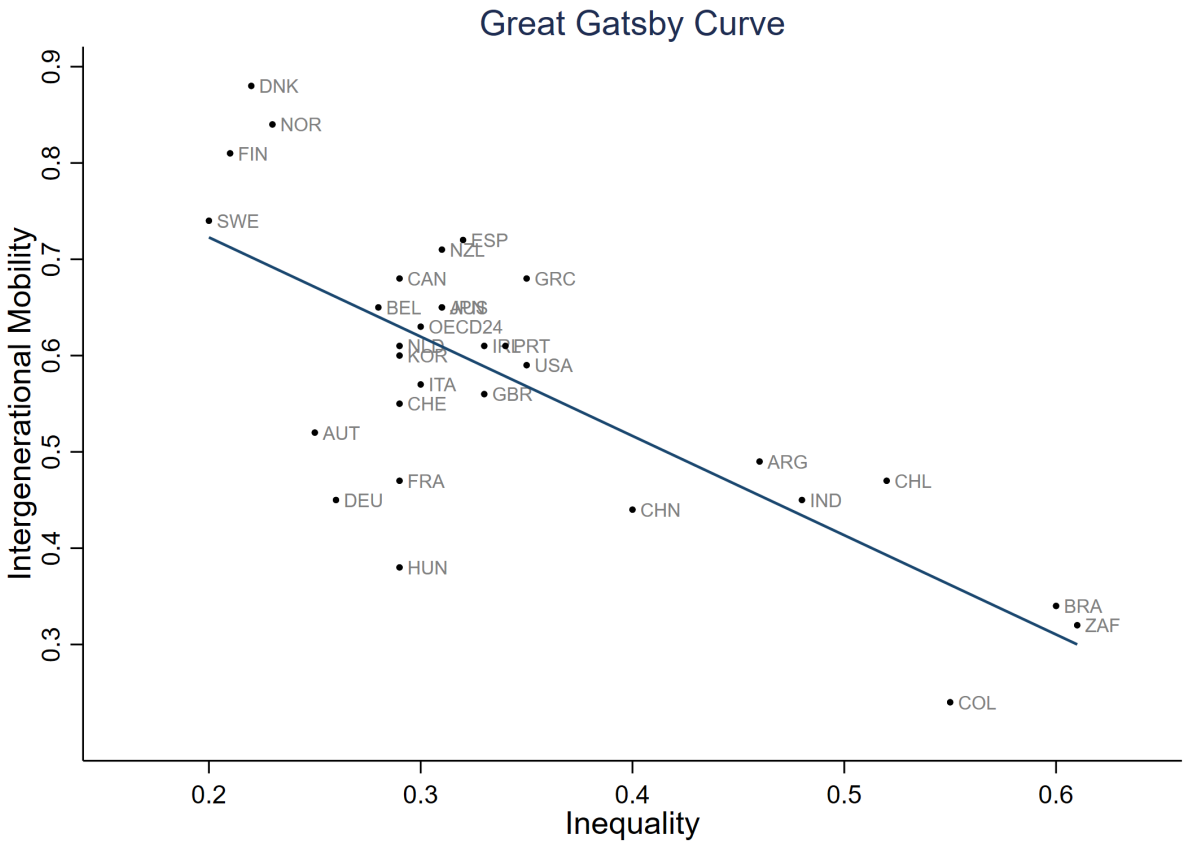


Figure 1: The Great Gatsby curve: inequality and intergenerational mobility

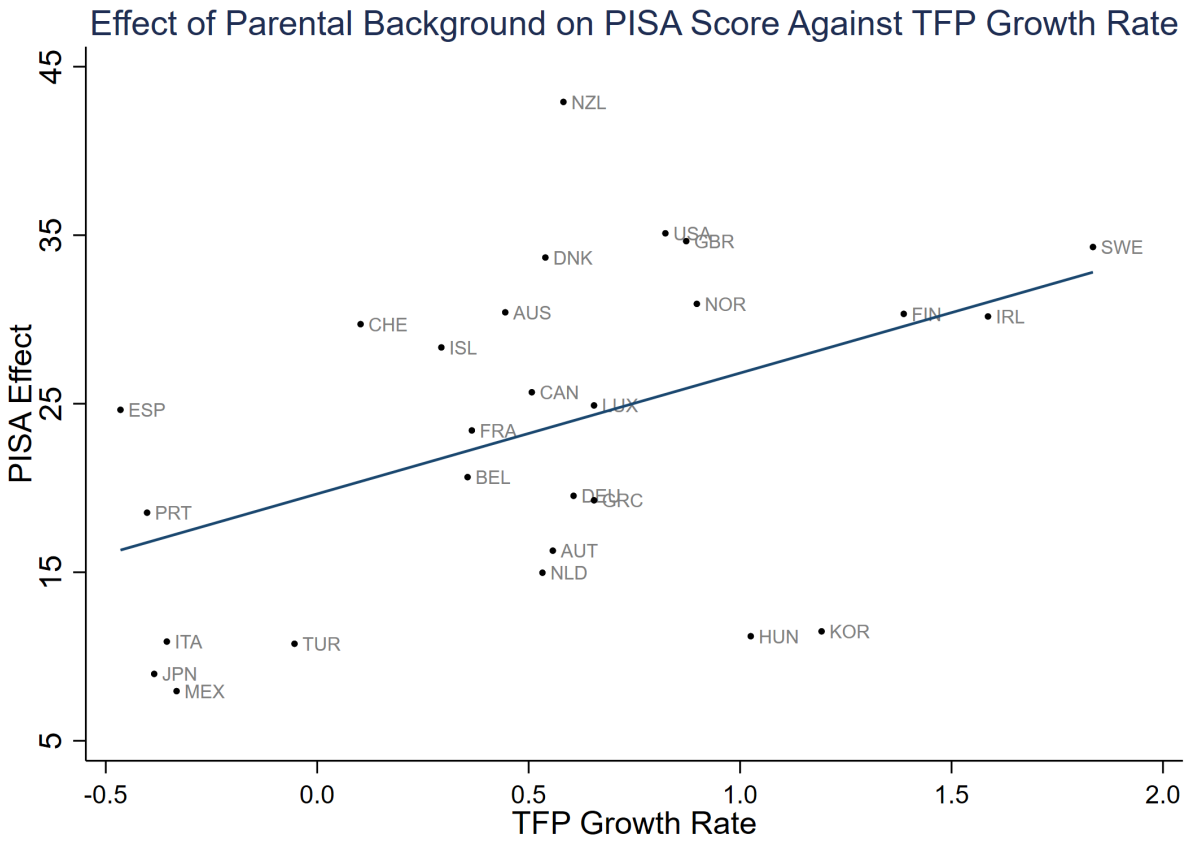


Figure 2: Productivity growth and the parental-background gradient in PISA scores

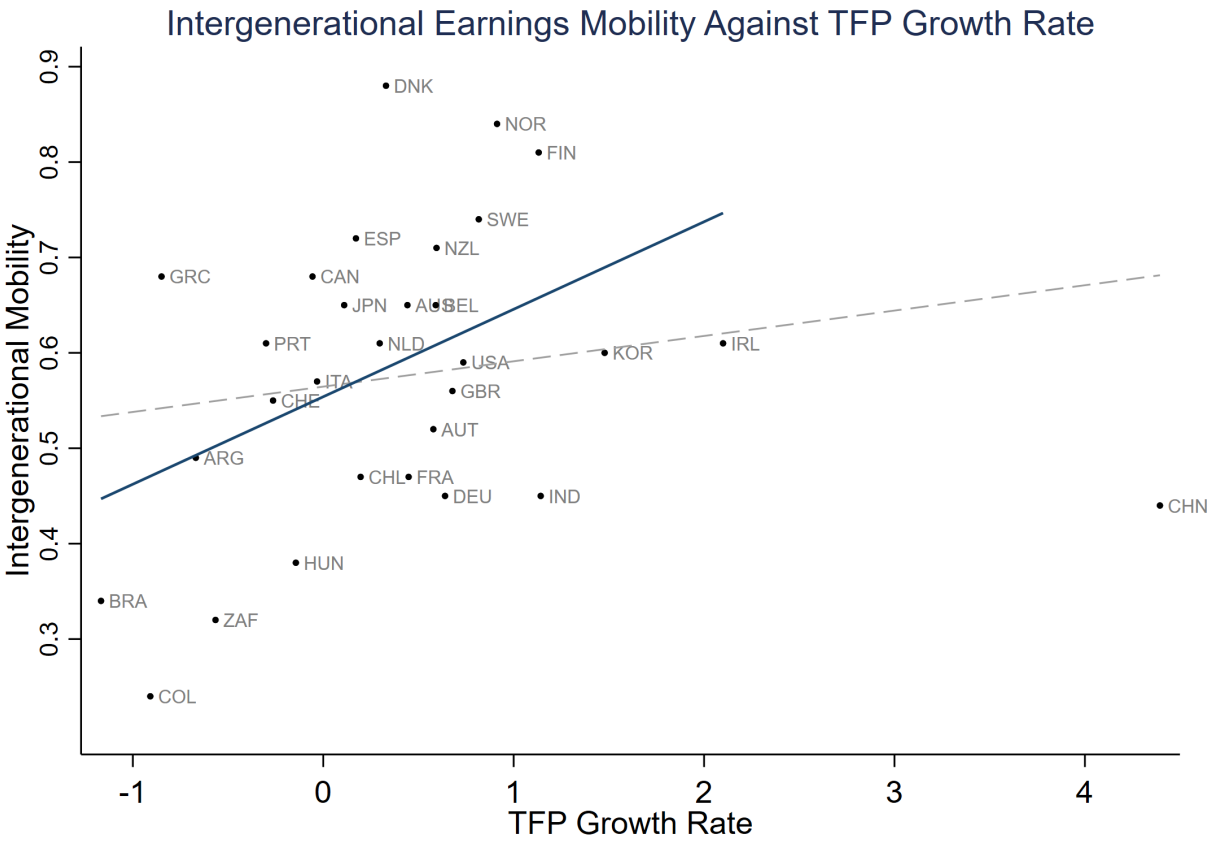


Figure 3: Childhood-aligned productivity growth and intergenerational mobility

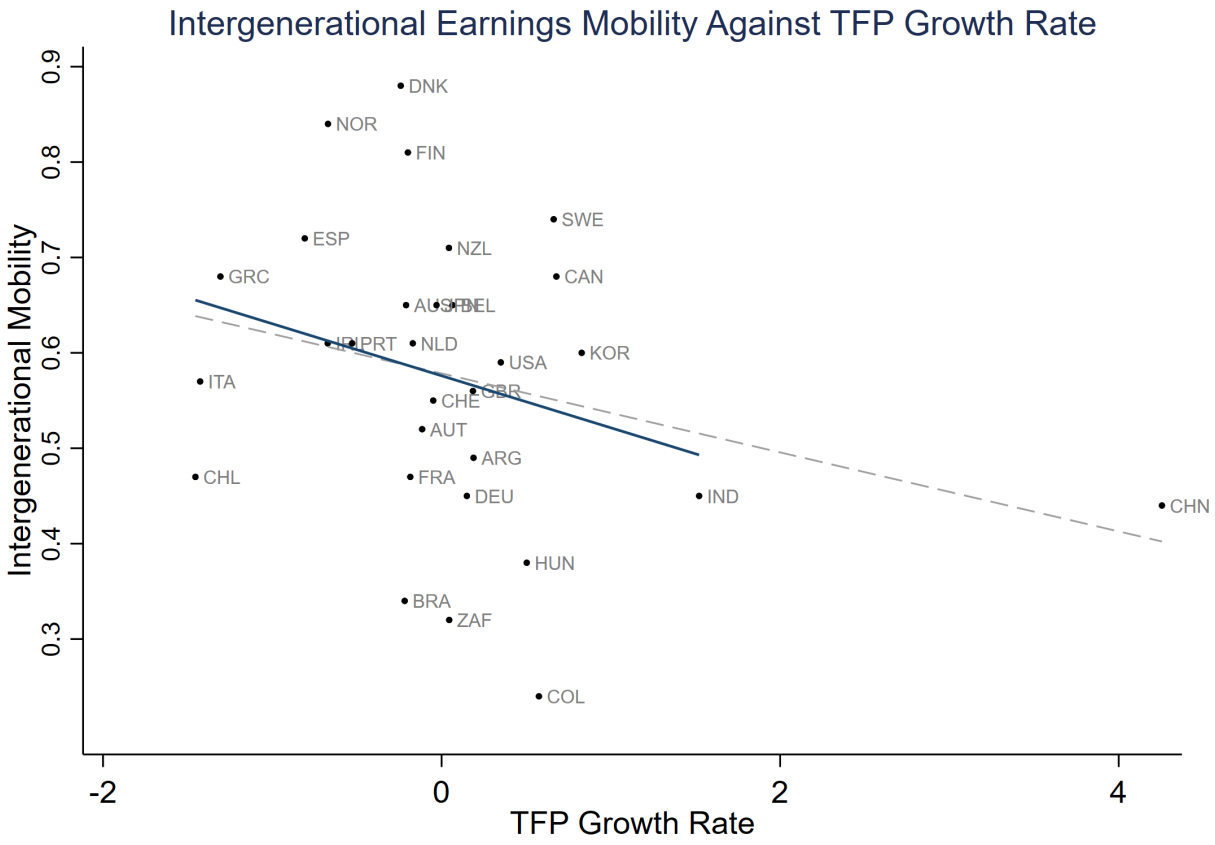


Figure 4: Adult-period productivity growth and intergenerational mobility

4 Theoretical Framework

Skills, tasks, and job offers. Each individual i in generation $g \in \{p, o\}$ enters the labor market with a vector of skills a_i^g. Parental resources, family investments, and inherited traits may all shape the joint distribution of parent and child skills. We take that distribution as given. The distinction is between skills as worker capacities and tasks as activities that firms require workers to perform (Autor et al., 2003; Acemoglu and Autor, 2011).

Firms offer jobs indexed by j. Each job specifies a bundle of tasks, with q_{jm}^g denoting the intensity of task m. The function $F_j^g(a)$ maps a worker's skills a to a vector of task productivities. For example, both reasoning and communication can contribute to solving a client's problem, and reasoning can also help with a different task, such as diagnosing a fault. If v_{jm}^g is the reward per unit of productivity in task m, the job offers log labor income

$$\ell_{ij}^g = b_j^g + \sum_m v_{jm}^g q_{jm}^g F_{jm}^g(a_i^g),$$

where b_j^g is a job intercept and F_{jm}^g is the component of the function for task m. In the linear special case, $F_j^g(a) = A_j^g a$, and the rows of A_j^g describe how each skill contributes to each task. More generally, the payoff to one skill can depend on the worker's other skills. Workers choose among available offers based on earnings and other job attributes. Technology can change which tasks firms require, how they reward those tasks, and how workers use their skills to perform them. These changes can affect both earnings within a job and the jobs workers choose.

From job assignment to income persistence. Let S_j^g summarize the skill requirements of the task bundle in job j. The empirical counterparts are occupational requirement indices. For the realized job j_i^g, write the linear projection of log income as

$$y_i^g = \alpha^g + (\tau^g)' S_{j_i^g}^g + u_i^g.$$

Here τ^g summarizes the income reward associated with entering jobs with different requirements, conditional on the controls used in estimation. It combines task rewards with the sorting of workers across jobs. The worker's endowment a_i^g and the job's requirements $S_{j_i^g}^g$ are distinct objects. Labor income is the direct outcome of the job offer model. The main estimates apply the same accounting to family income, which also reflects the earnings of other household members and nonlabor resources.

For linked parents and children, the regression of child log income on parental log income

yields the IGE. Suppressing controls for exposition, its population slope is

$$\beta = \frac{\mathrm{Cov}(y_i^o, y_i^p)}{\mathrm{Var}(y_i^p)} = (\tau^o)' \frac{\mathrm{Cov}(S^o_{j_i^o}, y_i^p)}{\mathrm{Var}(y_i^p)} + \frac{\mathrm{Cov}(u_i^o, y_i^p)}{\mathrm{Var}(y_i^p)}.$$

The fixed-profile decomposition holds occupational requirements constant and separates their income rewards from sorting by parental income. A rise in the reward to cognitive requirements can raise persistence when children from higher-income families enter more cognitively demanding jobs. A weaker parental-income gradient in those requirements can offset that increase.

Changing requirements and occupational mobility. Allowing requirements to vary across generations gives the accounting identity

$$(\tau^o)' S^o_{j_i^o} = (\tau^p)' S^p_{j_i^p} + (\tau^o - \tau^p)' S^p_{j_i^p} + (\tau^o)' (S^o_{j_i^o} - S^p_{j_i^p}).$$

The first two terms value the parent's job requirements at parental returns and then at the change in returns. The last term values the difference in requirements at child returns. Among pairs in the same occupation, that difference captures changes in measured job content. Among pairs in different occupations, it reflects the change in occupational assignment as well as any change in the requirements of those occupations. The counterfactual in Section 6 evaluates both occupations using the child-period profile to isolate the latter assignment margin. Each term contributes to the IGE through its association with parental income. Thus even substantial changes in job content can account for little persistence if their income value varies little with parental income.

5 Data and Occupational Skill Requirements

5.1 Intergenerational income observations

The analysis links adult children in the Panel Study of Income Dynamics (PSID) to their parents and combines these links with measures of occupational requirements. The estimation unit is a child–income-year observation. Children are observed at ages 30–35 and grouped by their recorded birth decade, from the 1950s through the 1990s. Both sons and daughters are included. Repeated observations on a child remain separate observations rather than being collapsed to a single income or occupation.

This pooled child-year design follows the PSID implementation of Nybom and Stuhler (2024), who combine annual child family income at ages 30–35 with parental log family income

averaged at child ages 15–17 and evaluate the IGE at age 33. Retaining eligible repeated observations uses more of the available panel information while clustering inference by child. Lee and Solon (2009) likewise use repeated annual observations, although over a broader child-age range and with different age controls. Our estimand is an age-standardized income association; neither pooling nor centering the interaction eliminates lifecycle bias in lifetime-income measurement. In particular, Nybom and Stuhler (2024) emphasize trends rather than an unbiased lifetime IGE level, and Nybom and Stuhler (2016) document remaining lifecycle concerns.[2]

Solon (1992) emphasizes the importance of measurement error and sample design for intergenerational income estimates, while Haider and Solon (2006) show that the relationship between current and lifetime earnings varies systematically over the life cycle. These concerns motivate the explicit income-age windows used here and caution in interpreting the youngest cohort. The resulting annual-income estimates are not interchangeable with lifetime-income elasticities.

The main outcome, Y_{it}, is the log real total family-unit money income attached to child i in income year t. Parental income, P_i, is the existing mean of eligible log real total family income recorded when the child is aged 15–17. It is an average of log income. We select the father when the original income and household-link criteria are satisfied and otherwise use the eligible mother. The household-link criterion requires a qualifying match at at least one date in the window; it does not require a match at every contributing date.

The starting sample contains 16,671 child-years and 4,738 distinct children with the income and control variables required by the original IGE specification. Missing occupations and incomplete requirement profiles reduce the samples used below. Table 1 distinguishes these samples explicitly.

For the 1990s birth-decade group, the available income observations cover only ages 30–32. The IGE specification evaluates the parental-income coefficient at age 33, so estimates for this group involve extrapolation. We report them separately as provisional observations rather than treating them as evidence of a precisely estimated recent trend. Some children also enter more than one birth-decade group because recorded birth-year assignments vary across their cleaned observations. Cohort-specific child counts therefore need not sum to the overall number of children, Appendix A records these limitations and the remaining sample checks.

[2]The auxiliary program that constructs a three-observation child-income average uses it for a dispersion statistic. It does not replace the annual income outcome in the IGE regression.

Table 1: Estimation samples by measurement design

Birth cohort	Fixed child-years	Fixed children	Three-source child-years	Three-source children	Dated child-years	Dated children	Transition child-years
1950s	4,933	1,084	4,771	1,032	4,771	1,032	4,456
1960s	2,992	952	2,877	903	2,819	896	2,638
1970s	2,619	1,042	2,422	964	2,247	948	2,157
1980s	2,872	1,207	2,625	1,118	2,386	1,024	2,494
1990s[†]	322	264	290	238	258	213	280

Notes: Fixed: broad O*NET 24.1 skills. Three-source: joint historical cognitive, noncognitive and manual model. Dated: cognitive-only common sample, also used for 22-family comparisons. Overall unique children are 4,445, 4,161, 4,024 and 3,921, respectively; child-years total 13,738, 12,985, 12,481 and 12,025. Cohort-specific child counts can overlap. The transition probabilities do not extrapolate to age 33; the income slopes do for the youngest group.

5.2 Occupations, measurement dates, and the common classification

PSID occupational codes change classification over time. Reporting waves through 2001 use the 1970 Census scheme, the 2003–2015 waves use the 2000 scheme, and waves from 2017 use the 2010 scheme. We map these codes to the harmonized IPUMS OCC 2010 classification, including its published category consolidations. We use published crosswalks and retain unmatched or ambiguous cases as missing. We do not match occupations by numerical proximity or infer an unobserved job history from occupation titles.

Among child observations with recorded occupation codes, 98.86% receive a common Census match. The corresponding rates for selected father and mother income-window records are 99.15% and 99.54%. A successful crosswalk is also not proof that the detailed job content represented by a common code is unchanged over time.

Income in the cleaned files refers to the calendar year before the reporting survey. We preserve that one-year correspondence when assigning occupational profiles. The interview occupation is thus a proxy for the occupation generating income in the preceding year; it does not establish that a respondent held the same job throughout that year. For parents, every available occupational record in the established income window is retained.

5.3 A fixed occupational requirement benchmark

The first measurement design assigns a fixed O*NET 24.1 profile, released in November 2019, to each common occupation. Cognitive requirements combine 21 cognitive abilities. Broad noncognitive requirements combine six Social Skills descriptors with 14 Work Styles

Table 2: Occupational requirement measures

Design	Cognitive	Noncognitive	Manual	Reference
Fixed O*NET 24.1	21 abilities	6 social + 14 styles	10 psychomotor + 9 physical	456 occupations
Historical DOT	3 aptitudes	5 temperaments	3 dexterity/coordination	267 occupations
Historical O*NET	7 abilities	5 social/styles	3 dexterity/coordination	423 / 448
Dated extension	3 DOT / 7 O*NET	Not included	Not included	Snapshot-specific

Notes: Each domain is represented by its oriented first principal component. Fixed O*NET uses Level for abilities and social skills and Importance for Work Styles. Historical cognition and manual abilities use Level in O*NET; the five historical noncognitive descriptors use Importance. DOT aptitudes are reversed ordinal categories and temperaments are presence indicators. Dated scores freeze DOT 1977 and O*NET 15.0 transformations within separate source blocks. Equal PCA reference weights differ from employment aggregation weights.

descriptors, excluding Analytical Thinking and Innovation to limit overlap with cognition. Manual requirements combine ten psychomotor and nine physical abilities. Table 2 summarizes the distinctions between this broad benchmark and the historical measures used later.

Ability and Social Skills inputs use Level ratings on the nominal 0–7 scale; Work Styles use Importance ratings on the 1–5 scale.

We first average usable O*NET specialty ratings equally within six-digit SOC groups, because specialty employment shares are unavailable. SOC means are then pooled within common OCC 2010 using fixed May 2018 national employment counts. In this fixed-profile benchmark, a group with an unsupported constituent employment count uses equal SOC weights for the whole group. Employment weights construct occupational profiles, they do not weight the PSID income regressions.

For each skill domain, we standardize the aggregated descriptors on the same reference of 456 occupations complete on all 60 inputs and retain the first principal component. Each reference occupation receives equal weight in PCA. We orient the component toward higher average requirements and normalize its score to reference mean zero and sample standard deviation one. The three domain scores are not orthogonalized against one another. PC 1 explains 62.1% of cognitive input variance, 54.3% of noncognitive variance, and 75.1% of manual variance. These are occupation-imputed requirements, not direct measurements of individual ability or personality.

The exact item inventories and scales are reported in Appendix A.4, Tables A.3–A.6. In particular, Level measures the amount or complexity required, whereas Importance measures how important the descriptor is for the occupation; the two scales are standardized as separate input columns rather than multiplied.

5.4 Historical profiles and dynamic cognitive requirements

The historical design uses the Dictionary of Occupational Titles (DOT) and archived O*NET releases. Its initial version assigns DOT 1977 to observations reported under the Census 1970 regime, O*NET 15.0 to Census 2000 reporting waves, and O*NET 24.1 to Census 2010 reporting waves. Each generation receives its own date-specific profile. The sources differ in descriptor content, so their indices measure related occupational constructs rather than an identical calibrated quantity of skill.

Historical cognition uses general learning, verbal aptitude, and numerical aptitude in DOT; O*NET uses deductive and inductive reasoning, oral and written comprehension, oral and written expression, and number facility. DOT aptitude categories run from higher to lower requirements and are reversed using six minus the original category. The historical noncognitive index uses five interpersonal and adaptability-related inputs; the manual index uses three dexterity and coordination inputs. These narrower constructs must not be confused with the broad fixed O*NET domains. Cognitive PC 1 explains approximately 86–90 percent of the historical cognitive descriptor variance. The initial three-source design has 267, 423, and 448 complete common occupation profiles, respectively.

DOT 1977 profiles use the original CPS-linked jobs and employment weights. The DOT 1991 sensitivity links revised ratings to that observed job universe through the documented conversion table; newly listed jobs without reference employment are not assigned invented weights. For O*NET, the initial three-source construction uses taxonomy-compatible employment references. A common group with a mapped SOC lacking a supported employment count is withheld. Known-weight but unrated cells remain in the coverage denominator. We permit descriptor means based on available rated constituents and retain coverage indicators rather than imposing an additional coverage threshold after inspecting outcomes.

The extended design focuses on cognition and assigns multiple dated profiles even within a Census coding regime. We construct DOT 1965, DOT 1977, and DOT 1991 profiles and eleven O*NET release profiles. Requirement assignment follows the information available by December 31 of the income year, with the dating calendar in Appendix Table A.1. O*NET 8.0 is retained for release-history diagnostics but is superseded by 9.0 at the relevant annual cutoff. Release dates are not the dates on which every underlying rating was collected, and repeated archived values do not count as fresh measurement.

In this extension, all DOT waves use the frozen DOT 1977 input means, standard deviations, PCA coefficients, and final score normalization. Every O*NET wave uses the frozen O*NET 15.0 transformation. Writing $X_{jk\ell}$ for descriptor k in common occupation j

and snapshot ℓ, the score within source block b is

$$S^b_{j\ell} = a_b + \sum_{k \in \mathcal{K}_b} c_{bk} X_{jk\ell}. \tag{3}$$

The constants a_b and c_{bk} are fixed across snapshots. Neither descriptors nor final scores are restandardized in each wave or child cohort. Common increases and changes in dispersion within a block can therefore remain visible. DOT and O*NET nevertheless retain different reference units; the common occupation code does not supply a cardinal bridge between them.

The extension also fixes the O*NET employment reference to May 2018 SOC employment across releases. Historical taxonomies are traced through official relations to the common classification, accepting a profile only when its supported lineage resolves to one common OCC 2010 target. This avoids duplication across common occupations. Available specialty composition can still change within that target. DOT 1965 requires separate reconstruction from job/function records and published job-to-Census weights, multiplied by reference Census employment before pooling. Its composition boundary with later DOT is retained as a qualification rather than labeled pure technological change.

6 Empirical Framework

6.1 Income persistence and the fixed-profile decomposition

For child cohort h, estimate

$$Y_{it} = \alpha_h + \beta_h P_i + \lambda_h P_i (A_{it} - 33) + X'_{it}\gamma_h + u_{it}, \tag{4}$$

where A_{it} is child age and X_{it} contains parental age and its square, child-age fixed effects, and child income-year fixed effects. The reported IGE, β_h, is evaluated at age 33. It is a conditional log-income slope, not a rank correlation, and it is not constrained by definition to lie between zero and one. We do not add education or other potential mediators to the approved control set.

Let S^o_{kit} denote the fixed-profile requirement for skill $k \in \{c, n, m\}$. The income projection and requirement projections are

$$Y_{it} = a_h + \sum_k \tau^o_{kh} S^o_{kit} + X'_{it} b_h + e_{it}, \tag{5}$$

$$S^o_{kit} = d_{kh} + \pi_{kh} P_i + \eta_{kh} P_i (A_{it} - 33) + X'_{it} q_{kh} + v_{kit}. \tag{6}$$

The income equation includes the skills jointly but excludes parental income and its interaction. We call τ^o_{kh} the return to skill, with the interpretation of a conditional income association. We call π_{kh} the intergenerational elasticity of skill for continuity with the decomposition terminology; because its outcome is a standardized level that can be negative, its exact units are requirement-score units per log point of parental income, a semi-elasticity.

Linearity of the common-sample projection yields

$$\beta_h = \sum_k \tau^o_{kh}\pi_{kh} + \beta^R_h, \tag{7}$$

where β^R_h is calculated by subtraction. We do not estimate a separate residual regression or impose zero persistence in the unexplained component. The identity partitions a conditional association; it does not identify causal mediation by measured individual skills.

6.2 Generation-specific returns and occupational change

For the historical designs, let $s^o_{it} = S_{j^o_{it},\ell^o_{it}}$ and $s^p_{ir} = S_{j^p_{ir},\ell^p_{ir}}$ be requirement vectors, with $\bar{s}^p_i = \sum_r w_{ir}s^p_{ir}$. Here $w_{ir} = 1/|\mathcal{R}_i|$ over the selected parent's matched income-window occupation records, so the weights sum to one without expanding the child-year regression sample. For each child cohort estimate

$$P_i = a^p_h + (\tau^p_h)'\bar{s}^p_i + Z_i'\delta^p_h + e^p_i, \tag{8}$$

$$Y_{it} = a^o_h + (\tau^o_h)'s^o_{it} + X_{it}'\delta^o_h + e^o_{it}. \tag{9}$$

The parent equation uses the same child-year observations as the child equation, preserving their repeated-observation weighting. Its controls, Z_i, contain parental age and its square and the shares of all original parental income-window records in each income year, with one reference year omitted. These shares are averaged year indicators, avoiding the assignment of a fictitious single year to a multi-year income average. The second-stage IGE controls remain X_{it} and the parental-income interaction in equation (4).

Define $\Delta\tau_h = \tau^o_h - \tau^p_h$ and $I_{itr} = \mathbf{1}\{j^o_{it} = j^p_{ir}\}$. Construct

$$B_{it} = (\tau^p_h)'\bar{s}^p_i, \qquad Q_{it} = (\Delta\tau_h)'\bar{s}^p_i, \tag{10}$$

$$W_{it} = \sum_r w_{ir}I_{itr}(\tau^o_h)'(s^o_{it} - s^p_{ir}), \tag{11}$$

$$T_{it} = \sum_r w_{ir}(1 - I_{itr})(\tau^o_h)'(s^o_{it} - s^p_{ir}). \tag{12}$$

The four terms represent baseline persistence, reevaluation, within-occupation change among

observed stayers, and between-occupation change among observed switchers. The skill differences are valued at *child* returns. This choice preserves the exact identity

$$B_{it} + Q_{it} + W_{it} + T_{it} = (\tau_h^o)' s_{it}^o. \tag{13}$$

Valuing the differences at parent returns instead would omit the interaction between return changes and requirement changes. Occupations are observed during the income measurement windows; the terms do not refer to unobserved occupations at labor-market entry.

Let $\mathcal{L}_h(V)$ denote the coefficient on P_i when outcome V is projected on precisely the regressors in equation (4). Applying this linear operator gives the five-part accounting

$$\beta_h = \underbrace{\mathcal{L}_h(B)}_{\beta_h^B} + \underbrace{\mathcal{L}_h(Q)}_{\beta_h^Q} + \underbrace{\mathcal{L}_h(W)}_{\beta_h^W} + \underbrace{\mathcal{L}_h(T)}_{\beta_h^T} + \beta_h^R. \tag{14}$$

The remainder again follows by subtraction. A negative between component means that income-valued child-minus-parent requirements have a negative conditional parental-income gradient. It does not mean that affluent children necessarily enter low-requirement jobs, nor that switching occupations causally improves opportunity.

6.3 Counterfactual requirements and aggregation

The stayer definition in equation (11) excludes requirement evolution in a switcher's parental occupation. A complementary accounting evaluates that occupation under the child's profile, $\tilde{s}_{itr}^p = S_{j_{ir}^p, \ell_{it}^o}$, and writes

$$W_{it}^{\text{cf}} = \sum_r w_{ir} (\tau_h^o)' (\tilde{s}_{itr}^p - s_{ir}^p), \tag{15}$$

$$T_{it}^{\text{cf}} = \sum_r w_{ir} (\tau_h^o)' (s_{it}^o - \tilde{s}_{itr}^p). \tag{16}$$

Their sum equals $W_{it} + T_{it}$ on common observations. The first term asks how requirements in the parental occupation changed, regardless of whether the child entered it; the second compares occupations using the same child-period profile. This is an accounting counterfactual, not a model of the occupation a child would choose under an intervention. Missing counterfactual profiles require a common-sample comparison rather than silently dropping parent records and renormalizing their weights.

We also use the 22 civilian SOC major occupational families. One check changes only the definition of occupational equality; another pools raw descriptors within families before applying the frozen scoring rule and reestimating returns. The second check changes the

explanatory variables, not just the within/between labels. It can alter baseline persistence, reevaluation, and total explained persistence even on identical child-years. It is this family-profile design that we emphasize for the dated cognitive robustness results.

6.4 Normalization, inference, and interpretation

As a sensitivity check, standardize each generation's log income within the child's cohort and exact estimation sample. Repeated parental income values retain their child-year weighting. Let $f_h = \mathrm{SD}(P)/\mathrm{SD}(Y)$. In the fixed-profile decomposition all contributions rescale by f_h. In the generation-specific accounting, however,

$$\beta_h^{B,z} = \beta_h^B, \qquad \beta_h^{Q,z} = f_h(\beta_h^B + \beta_h^Q) - \beta_h^B, \tag{17}$$

while within, between, remainder, and IGE rescale by f_h. This distinction follows because the parent return and parental-income gradient rescale in opposite directions. We reestimate or verify these identities directly rather than applying a common factor to all five components.

Uncertainty for products, generated components, and remainders uses 999 whole-child bootstrap draws within cohort. A draw retains all observations of each sampled child, reestimates both return equations and all component projections, and recomputes income standardization where relevant. Paired specification comparisons use identical child draws. This procedure preserves covariance among estimated returns and gradients; bootstrapping is needed here for inference on their products and accounting differences, not because estimating returns itself establishes nonclassical measurement error. Individual coefficients in the fixed-profile trend figure use child-clustered sandwich intervals; the tables report bootstrap percentile intervals. Measurement profiles, PCA coefficients, and crosswalks remain fixed, so these intervals condition on the occupational measurement system. They also do not add a separate sibling-level cluster.

7 Fixed Requirements: Returns and Intergenerational Sorting

The fixed-profile benchmark isolates sorting across a stable occupational requirement system. Cognitive requirements have the largest positive measured contribution in every cohort. Table 3 and Figure 5 distinguish the income return from the intergenerational requirement gradient. The cognitive return is 0.212 for the 1950s cohort, falls to 0.168 for the 1960s, and reaches 0.351 for the 1980s. Over the 1950s–1980s comparison, the cognitive parental-income

Table 3: Fixed O*NET requirements and total-income persistence

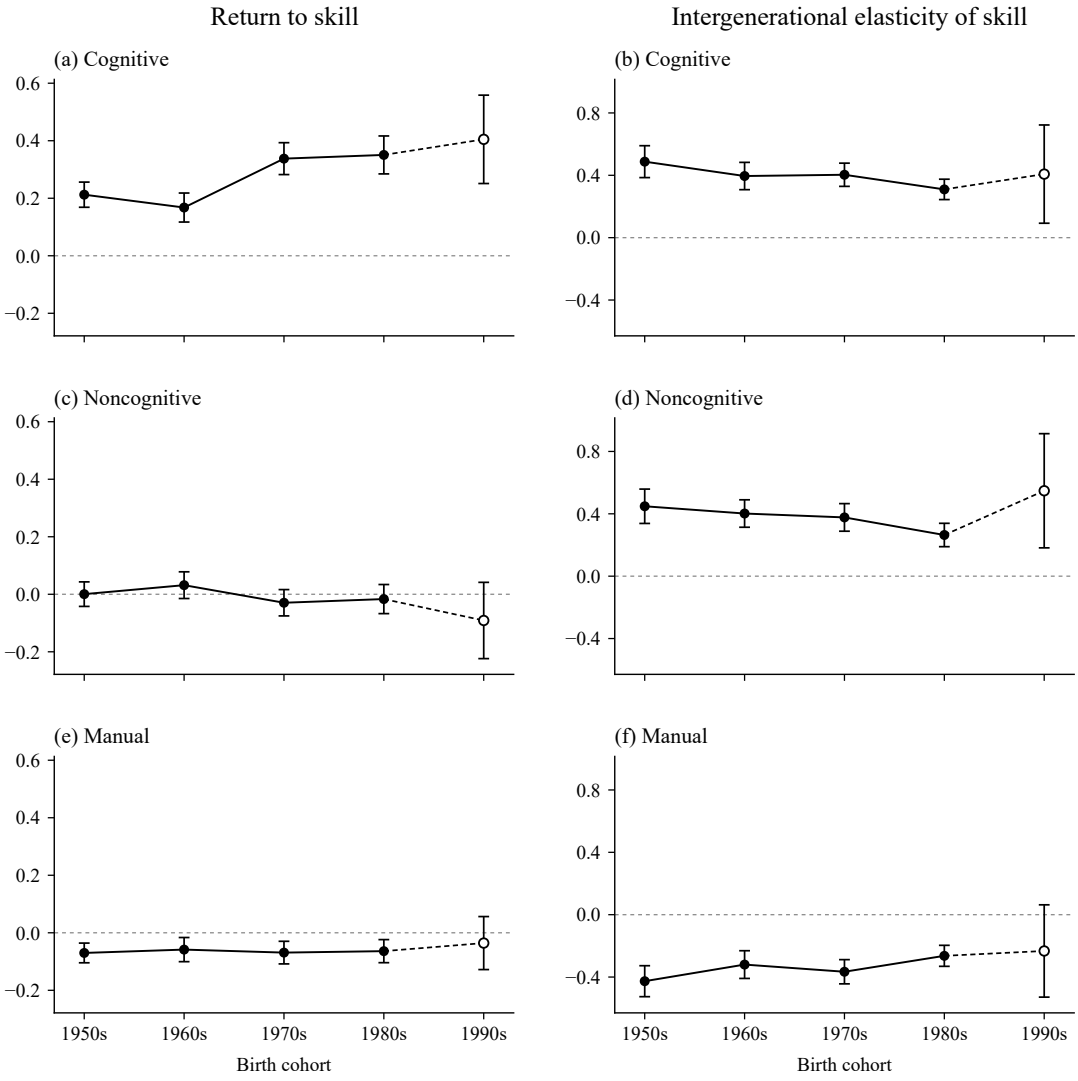
Cohort	IGE	τ_c^o	π_c	$\tau_c^o\pi_c$	$\tau_n^o\pi_n$	$\tau_m^o\pi_m$
1950s	0.363 [0.288, 0.435]	0.212 [0.170, 0.258]	0.488 [0.392, 0.586]	0.104 [0.073, 0.136]	0.000 [-0.019, 0.021]	0.030 [0.014, 0.049]
1960s	0.374 [0.312, 0.445]	0.168 [0.121, 0.218]	0.395 [0.308, 0.478]	0.066 [0.044, 0.093]	0.013 [-0.006, 0.030]	0.019 [0.004, 0.038]
1970s	0.391 [0.319, 0.461]	0.338 [0.285, 0.390]	0.403 [0.333, 0.478]	0.136 [0.103, 0.171]	-0.011 [-0.028, 0.006]	0.025 [0.010, 0.042]
1980s	0.347 [0.292, 0.409]	0.351 [0.284, 0.419]	0.310 [0.243, 0.379]	0.109 [0.078, 0.146]	-0.004 [-0.018, 0.008]	0.017 [0.006, 0.030]
1990s†	0.506 [0.182, 0.861]	0.405 [0.256, 0.564]	0.408 [0.103, 0.732]	0.165 [0.042, 0.334]	-0.050 [-0.144, 0.022]	0.008 [-0.015, 0.042]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. The three skill returns enter the income equation jointly. Requirements retain their occupation-reference normalization. The gradient π_c is evaluated at age 33 and is a standardized-level-on-log-income slope.

gradient falls from 0.488 to 0.310. These offsetting point-estimate changes leave the cognitive product relatively similar at the two endpoints, 0.104 and 0.109, even though the intermediate cohorts differ.

This pattern illustrates why a rising association between income and cognitive-intensive occupations need not generate a parallel rise in IGE. The contribution depends jointly on the return and the extent to which access to such occupations is differentiated by parental income. It is an interpretation of the accounting, not a formal joint test that a price trend is offset by a supply trend. Fixed occupation profiles cannot establish changes in requirements within a job, and imputed requirements cannot distinguish acquired individual skill from occupational access, job allocation, or household income composition.

Broad noncognitive returns are small and imprecisely estimated conditional on the other domains. Manual returns are negative, while parental income predicts lower manual requirements, producing positive manual products. These signs refer to occupational associations; they do not imply that reducing a person's manual ability would raise income. The three skill products account for approximately 24–38 percent of IGE across cohorts, with a substantial remainder. Alternative narrow indices and separate social, character, psychomotor, and physical domains produce similar combined contributions but redistribute individual components (Appendix B). Thus a cognitive focus is a parsimonious presentation of the largest component, not a restriction that other domains are known to be irrelevant.



Notes: Bars show 95% child-clustered confidence intervals. Returns are log family income per skill SD.
Elasticities are skill SDs per log point of parental income, evaluated at child age 33.
Hollow markers and dashed segments identify the 1990s age-33 extrapolation (observed ages 30–32).

Figure 5: Fixed occupational requirements: returns and intergenerational gradients

Notes: Broad O*NET 24.1 indices;13,738 child-years. Left: joint income coefficients; right: requirement-on-parental-income gradients at age 33. Bars are child-clustered 95 percent t intervals, not the bootstrap intervals in Table 3. Hollow youngest-cohort markers indicate extrapolation. Each column has distinct units.

8 Historical Requirements and the Sources of Income Persistence

8.1 The three-snapshot benchmark

Introducing generation-specific historical profiles makes the distinction between inherited occupational position and subsequent reevaluation explicit. The joint three-domain model uses 12,985 child-years and 4,161 children. Its total-income IGE ranges from 0.338 to 0.431 through the 1980s cohorts. Baseline persistence is positive, and between-occupation contributions are negative in each cohort. Cognitive net contributions are 0.112, 0.082, 0.144, and 0.126 for the 1950s–1980s. Older cohorts retain some positive noncognitive contribution, while manual components can offset positive cognitive components. Appendix Table B.2 retains the full three-domain accounting.

Within-occupation contributions in this initial design are small, but the interpretation is constrained by its measurement calendar. In particular, the 1950s within term is exactly zero when both generations receive the same DOT 1977 profile. Only approximately 2–6 percent of matched child–parent occupation exposures have the same detailed code. A small stayer component can therefore reflect both limited same-occupation exposure and infrequent remeasurement. The counterfactual in equation (16) addresses the first limitation; the dated extension addresses the second.

8.2 Multiple dated cognitive profiles

The dated extension includes cognition alone in each return equation; its coefficients are reestimated rather than taken from the joint three-domain model. Dated and three-snapshot measurement variants are compared on 12,481 identical child-years and 4,024 children. Eligibility requires a child score and a score for every parent record originally used in that child's matched requirement window. We do not improve coverage by changing the selected parent or dropping difficult parent dates. This retains 96.1 percent of the three-snapshot matched child-years.

Table 4 reports the resulting five-part decomposition. Baseline persistence is positive in every cohort. Reevaluation is not uniformly positive: it is 0.019, -0.051, -0.023, 0.048, and 0.159. Child cognitive returns are 0.245, 0.227, 0.259, 0.371, and 0.357. Thus the data show a higher return in the 1980s than the oldest cohorts, but not a monotonic increase in every birth decade. The latest cohort remains too imprecise for a strong trend conclusion.

The within-occupation contribution is no longer mechanically zero for the oldest cohort, because parents and children can receive distinct DOT snapshots even within the same Census

Table 4: Dated cognitive requirements: total-income IGE decomposition

Cohort	IGE	Baseline	Reevaluation	Within	Between	Remainder
1950s	0.365 [0.286, 0.441]	0.197 [0.157, 0.241]	0.019 [-0.021, 0.055]	-0.0028 [-0.0061, 0.0005]	-0.086 [-0.122, -0.055]	0.239 [0.166, 0.310]
1960s	0.367 [0.299, 0.444]	0.240 [0.196, 0.288]	-0.051 [-0.088, -0.012]	0.0005 [-0.0007, 0.0022]	-0.097 [-0.125, -0.071]	0.275 [0.213, 0.343]
1970s	0.429 [0.354, 0.513]	0.173 [0.126, 0.224]	-0.023 [-0.061, 0.010]	-0.0003 [-0.0029, 0.0021]	-0.033 [-0.065, 0.002]	0.313 [0.249, 0.383]
1980s	0.323 [0.257, 0.383]	0.117 [0.079, 0.158]	0.048 [0.021, 0.072]	-0.0020 [-0.0056, 0.0006]	-0.041 [-0.089, 0.001]	0.201 [0.145, 0.256]
1990s†	0.613 [0.159, 1.029]	0.142 [0.050, 0.263]	0.159 [0.034, 0.317]	0.0030 [-0.0034, 0.0114]	-0.199 [-0.395, -0.017]	0.509 [0.061, 0.927]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Components are separately projected on parental income with the same IGE controls; the remainder is obtained by subtraction. Rounding can affect displayed add-up.

coding regime. Nevertheless, detailed-occupation within estimates range only from -0.0028 to 0.0030, with all five percentile intervals including zero. Between-occupation estimates are negative throughout, with intervals excluding zero for the 1950s, 1960s, and provisional 1990s cohorts. For the 1980s, baseline persistence of 0.117 is partly offset by a between contribution of -0.041 and supplemented by reevaluation of 0.048, leaving a net cognitive contribution of approximately 0.122. The remainder is approximately 0.201.

These quantities clarify the empirical content of the negative between component. Children from lower-income backgrounds can make larger gains relative to lower parental occupational starting points while remaining less likely to reach the highest-requirement destinations. The regression compares income-valued changes against parental income, conditional on the common controls. It does not directly compare destination levels or identify the causal benefit of moving occupations. Section 9 examines starting positions and destinations separately.

8.3 Broader occupational families

The 22-family profile analysis uses the same 12,481 child-years, unchanged income variables and parent-record weights, and the same dated source calendar. We aggregate raw descriptors to families before applying the frozen scoring coefficients. All 22 families are observed in each of the fourteen constructed snapshots. Family return coefficients are reestimated; the exercise is not a mechanical relabeling of detailed-occupation terms.

The within-family contribution remains modest in magnitude, with an absolute point-estimate share of at most 1.6 percent of IGE. Statistical insignificance does not hold universally: the 1950s estimate is -0.0038, with interval [-0.0069,-0.0008], and the 1970s estimate is 0.0069, with interval [0.0033,0.0113]. The paired family-minus-detailed within contrast also

excludes zero in the 1970s. Between-family coefficients are less negative than their detailed counterparts. Appendix Table B.5 and Figure B.3 display the comparison. The evidence supports a small measured within contribution under coarser occupational definitions, while allowing economically distinct and statistically detectable changes in particular cohorts.

8.4 Do requirements rise while occupational ranks remain stable?

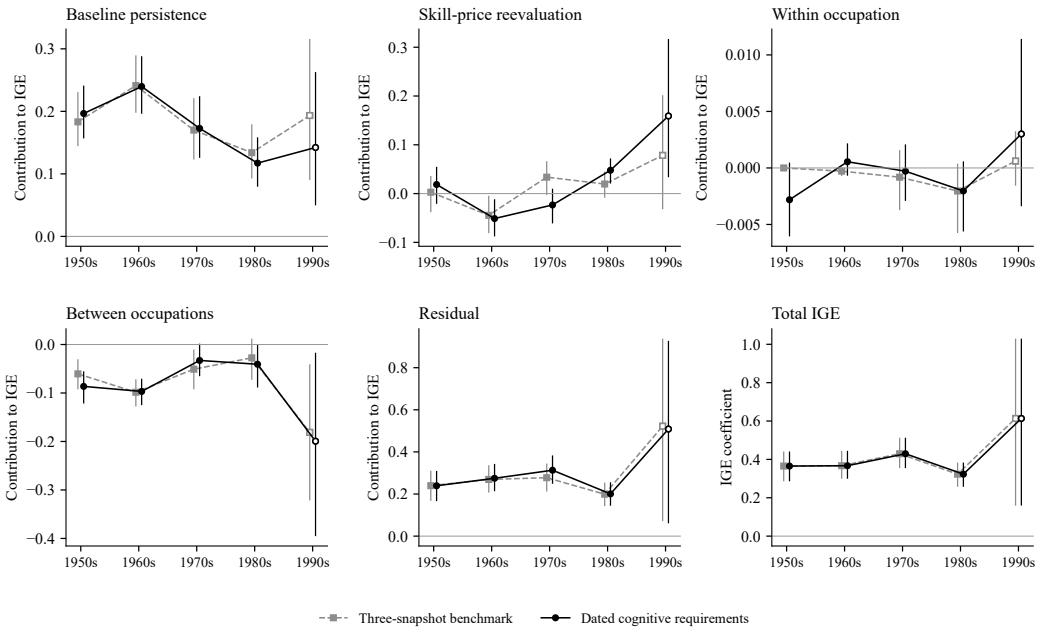
We next describe the requirement profiles themselves, independently of the PSID income regressions. The main rank and distribution comparisons hold the detailed occupation set fixed within each source block: 267 groups across DOT 1965–1991 and 392 across all eleven O*NET releases. A later O*NET 15.0–27.1 comparison contains 425 balanced groups. Equal occupation weights describe the reference occupation distribution; fixed employment weights provide a sensitivity check. The 22-family series is reported alongside a sensitivity that also fixes its detailed constituent occupation set before raw-descriptor aggregation.

Occupational hierarchies are persistent. Endpoint Spearman correlations are 0.930 for detailed DOT occupations and 0.886 for detailed O*NET occupations, rising to 0.983 over the later O*NET period. Family correlations are 0.964, 0.941, and 0.988. Kendall correlations confirm substantial, but imperfect, ordering persistence. These are descriptive comparisons of archived profiles, without treating linked occupations or repeated ratings as independent observations for conventional significance tests.

The level evidence is more qualified than a general claim of secular upgrading. Mean DOT cognition changes little and slightly declines in this aptitude index. O*NET 5.1–27.1 shows a mean increase of 0.722 detailed-occupation reference units, with approximately 80 percent of occupations increasing. Yet much of the movement occurs in early releases as the fixed-weight share of legacy analyst ratings falls from 84.6 percent to zero by release 15.0. We cannot separate actual requirement change from replacement of legacy ratings, revised anchors, and changes in covered specialty composition.

Over O*NET 15.0–27.1, the equal-occupation mean increase is only 0.011 units for detailed occupations and 0.045 for families; balanced family constituents reduce the latter to 0.029. Dispersion is nearly unchanged in this later period. Across the entire O*NET window, detailed SD actually falls from 1.676 to 1.000 rather than widening. Table 5 and Figure 7 show these differences. Our measured profiles therefore support persistent hierarchies and limited later-period growth, not a uniformly increasing or fanning-out cognitive distribution over the full historical period.

Stable ordering is compatible with nonzero requirement changes, and it is not equivalent to fixed standardized scores. Conversely, small within contributions to IGE do not prove



Notes: Cognitive only; same child-year sample and original parental-record weights. Bars: 95% child-cluster bootstrap intervals, 999 draws. Hollow markers: 1990s cohort, whose age-33 coefficients extrapolate beyond observed ages 30–32. Source units are reference-normalized. Panels have separate vertical scales. Residual is obtained by subtraction; no residual equation is estimated.

Figure 6: Dated cognitive requirements and total-income persistence

Notes: Cognitive-only models;12,481 common child-years. Black: dated requirements; gray: the original three-snapshot benchmark reestimated with cognition alone on the same sample. Five contributions sum to IGE. Bars are 95 percent whole-child bootstrap intervals. The youngest-cohort age 33 target is outside its observed ages 30–32. Panels have separate vertical scales.

Table 5: Requirement levels, dispersion, and occupational rank stability

Period	Unit	N	Spearman	Kendall	Δ mean	Rising (%)	Initial SD	Final SD
1965–1991	Detailed	267	0.930	0.770	-0.036	48.3	1.057	1.000
1965–1991	22 families	22	0.964	0.879	-0.024	36.4	1.058	1.084
2003–2022	Detailed	392	0.886	0.697	0.722	79.8	1.676	1.000
2003–2022	22 families	22	0.941	0.810	0.317	81.8	1.476	0.945
2010–2022	Detailed	425	0.983	0.886	0.011	51.5	1.009	1.014
2010–2022	22 families	22	0.988	0.931	0.045	72.7	0.938	0.945

Notes: Balanced occupational support within each block; equal occupation or family weights. Later O*NET starts with 15.0; full O*NET starts with 5.1. DOT and O*NET have different reference units. SD is a descriptive population moment. Rank coefficients are unweighted. A rise in the score is not a percentage increase in cognitive ability. Crosswalk and rating changes qualify temporal interpretation; no iid occupation-level confidence intervals are asserted.

that requirements were unchanged: an IGE component depends on the conditional parental-income gradient of an income-valued change. The diagnostics here do not estimate that additional relationship. Nor do related results on task change within occupations, such as Autor et al. (2003), imply a particular contribution to intergenerational income covariance. Our aptitude-based cognitive index and income-persistence estimand differ from their task measures and outcomes.

9 Occupational Movement and Access by Family Background

9.1 Two shared occupational hierarchies

We complement the income accounting with two kinds of matrices. Occupational transition matrices describe destinations conditional on the parent's occupational group. Access matrices describe destinations conditional on parental income. The first asks whether children move relative to their occupational origins; the second asks who reaches cognitively demanding occupations.

We define five cognitive groups using quintile cutoffs in a frozen equal-occupation reference. These are quintiles of reference occupations, not of workers. Ties remain in the same group and can prevent exact twenty-percent group sizes. In the fixed hierarchy, both generations' occupations receive the 2019 seven-input cognitive profile used in the historical construction. This is distinct from the broad 21-input fixed O*NET measure in Section 7. In the child-period hierarchy, both occupations receive the three-snapshot profile assigned to the child's observation. The latter evaluates the parental occupation under the opportunities represented by that child-period profile, as in the counterfactual between term. It does not compare each generation under a different hierarchy.

The common matrix sample contains 12,025 child-years and 3,921 children. All originally contributing parental records must be scoreable under both shared profiles; we do not drop individual parent records and renormalize their weights. Let g^p_{itr} and g^o_{it} denote the resulting occupational groups. The occupational transition probability is

$$M_{ab,h} = \frac{\sum_{it \in h} \sum_r w_{ir} \mathbf{1}\{g^p_{itr} = a, g^o_{it} = b\}}{\sum_{it \in h} \sum_r w_{ir} \mathbf{1}\{g^p_{itr} = a\}}. \tag{18}$$

Rows sum to one when the denominator is positive. The diagonal means the same cognitive group, not the same occupation. Income-access probabilities replace the origin condition with the cohort-specific quintile of parental income and give each child-year total weight one.

Table 6: Occupational movement and parental starting positions

Hierarchy	Cohort	Upward (%)	Origin-adjusted (%)	Same group (%)	Same occupation (%)
Fixed	1950s	38.8	35.4	29.3	5.9
Fixed	1960s	40.8	38.7	29.3	4.9
Fixed	1970s	32.7	34.5	26.5	4.0
Fixed	1980s	28.5	35.8	28.8	2.1
Child-period	1950s	42.1	35.5	28.3	5.9
Child-period	1960s	45.5	38.1	27.2	4.9
Child-period	1970s	37.1	34.9	28.6	4.0
Child-period	1980s	29.7	36.0	28.9	2.1

Notes: The common sample contains 12,025 child-years and 3,921 children. Origin-adjusted upward movement fixes the pooled 1950s–1980s fixed-hierarchy origin shares. It does not fix destination margins. Parent records receive original within-window weights, summing to one per child-year. Same group means a cognitive requirement group, not the same OCC 2010 occupation. Youngest-cohort descriptive matrices are retained in the source results, without age 33 extrapolation.

9.2 Starting positions and upward movement

Under the child-period hierarchy, raw upward-group movement falls from 42.1 percent in the 1950s cohort to 29.7 percent in the 1980s. Under the fixed hierarchy, it falls from 38.8 to 28.5 percent. These comparisons alone do not establish worsening conditional mobility: parents near the top have fewer upward destinations available.

We therefore standardize upward probabilities to one pooled fixed-hierarchy origin distribution for the 1950s–1980s cohorts. If ω_a is this reference share and U_{ah} the probability of upward movement conditional on origin a, the standardized rate is $\sum_a \omega_a U_{ah}$. For the child-period hierarchy, the 1950s and 1980s rates become 35.5 and 36.0 percent; the corresponding fixed-hierarchy rates are 35.4 and 35.8 percent. Much of the raw endpoint decline is associated with changes in parental starting positions under this accounting. This adjustment holds origin shares fixed, not destination margins, so it does not establish unchanged origin–destination association. Table 6 reports the comparison.

9.3 Access to high-cognitive occupations

Large differences by parental income coexist with offsetting between-occupation IGE components. In the 1980s child-period hierarchy, children in the lowest parental-income quintile reach the highest cognitive group with probability 11.8 percent, compared with 45.3 percent in the highest-income quintile. A larger sample requiring child profiles but no parental occupation yields a similar gap, 34.9 percentage points rather than 33.5. This suggests that

observing parental occupations alone does not generate the access gradient in the available sample, although broader selection into observed income and employment remains.

Even within some parental cognitive groups, income-related destination differences are substantial. For the 1980s child-period hierarchy, the high-minus-low-income gap in top-group access is 32.5 percentage points for origin group 4, with interval [16.8,46.8], and 30.3 points for origin group 5, with interval [17.1,41.7]. Other cells are sparse; origin group 1 has only four high-income children. We therefore do not infer a universal pattern from every large conditional point estimate. Appendix Figure C.2 retains the conditional-access display and its support qualifications.

The matrices describe unadjusted occupational positions, while the negative between component is a conditional slope of an income-valued difference. A child can gain more relative to a low parental origin yet remain less likely to reach the top than a child from an affluent family. These observations reconcile negative between contributions with unequal destination access without equating occupational requirement movement to changes in tested cognitive ability, wages, or welfare.

10 Sensitivity and Scope of the Empirical Evidence

10.1 Alternative indices, parental selection, and counterfactuals

The fixed-profile analysis retains a narrow cognitive/self-regulation specification and a five-domain specification separating character, social, psychomotor, and physical requirements. Their combined skill contributions are close to those in the broad benchmark, although conditional subdomain signs differ. The historical three-domain analysis additionally retains father-selected pairs, standardized log incomes, DOT 1991 historical ratings, and the parental-occupation counterfactual. Father-only means fathers paired with eligible children of either sex, not a literal father–son restriction. It changes sample composition as well as parental coverage.

The counterfactual comparison is conducted on observations with the additional parental-occupation profile. Its cognitive within estimates remain modest and imprecise, including -0.0026 for the 1970s and -0.0014 for the 1980s. These exercises preserve the interpretation that the largest measured requirement-related component is cognitive while avoiding a claim that all other skill channels are zero. Noncognitive two-PC combinations and equal-input means were prepared as measurement alternatives but have not been substituted into these decomposition regressions. Balanced cognitive-facet and three-PC specifications are not treated as estimated robustness results.

10.2 Labor income in both generations

An appendix extension repeats the three-snapshot joint-skill procedure using individual labor income for both generations. Child earnings follow the child's household relationship rather than automatically using household-head earnings. Parent selection follows labor-income availability, and occupational exposure is rebuilt for that selected parent's earnings window. Thus the exercise does not regress child earnings on parental total income and label the result labor-income IGE.

The labor sample contains 12,562 child-years and 4,065 children. Through the 1980s cohorts, labor-income IGE lies between 0.181 and 0.217; cognitive net contributions are 0.079, 0.098, 0.136, and 0.086. The relative importance of cognition persists, but the total-income pattern is not simply a common labor-market price trend. In common-child-year comparisons, total-income IGE exceeds labor-income IGE in the 1950s–1980s, with paired intervals excluding zero. Appendix Table D.2 reports these comparisons separately from the extended cognitive sample.

The total outcome is family-unit money income whereas the labor outcome is individual earnings. Their difference therefore reflects recipient units as well as income components; it cannot be assigned solely to capital income, transfers, or another nonlabor source. Annual earnings also combine pay, hours, and employment attachment. We preserve the original age controls in this extension, including the previously documented discrepancy when the labor-selected parent or earnings window differs from the original total-income parent window. Appendix D records this unresolved alignment issue instead of treating it as a completed correction.

Taken together, these exercises describe how occupational requirements and their income associations account for observed persistence under alternative measurement choices. They do not identify causal technological shocks, individual skill transmission, or a structural occupational-choice model. The historical source boundary, occupation aggregation, incomplete parent windows, and provisional youngest cohort remain central limits on the interpretation.

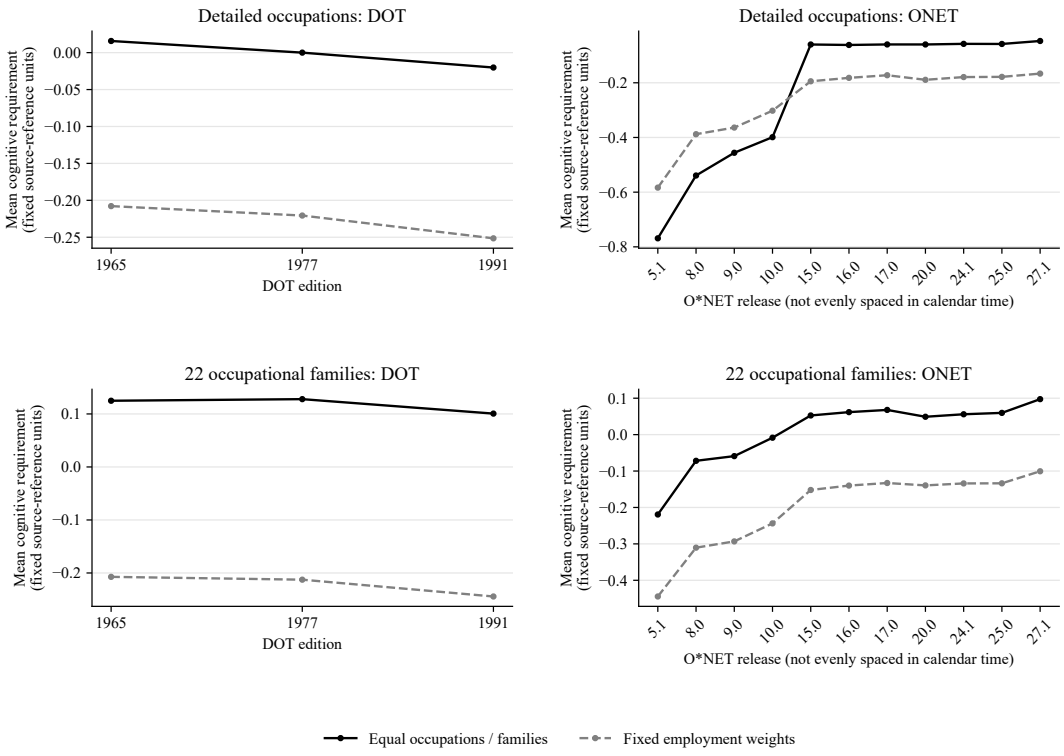


Figure 7: Mean cognitive requirements within source blocks

Notes: Balanced reference occupation panels; equal occupation weights and fixed employment weights. No wave restandardization. DOT and O*NET have different reference units, and releases are not uniformly spaced in calendar time. Early rating replacement and changing specialty coverage qualify level changes.

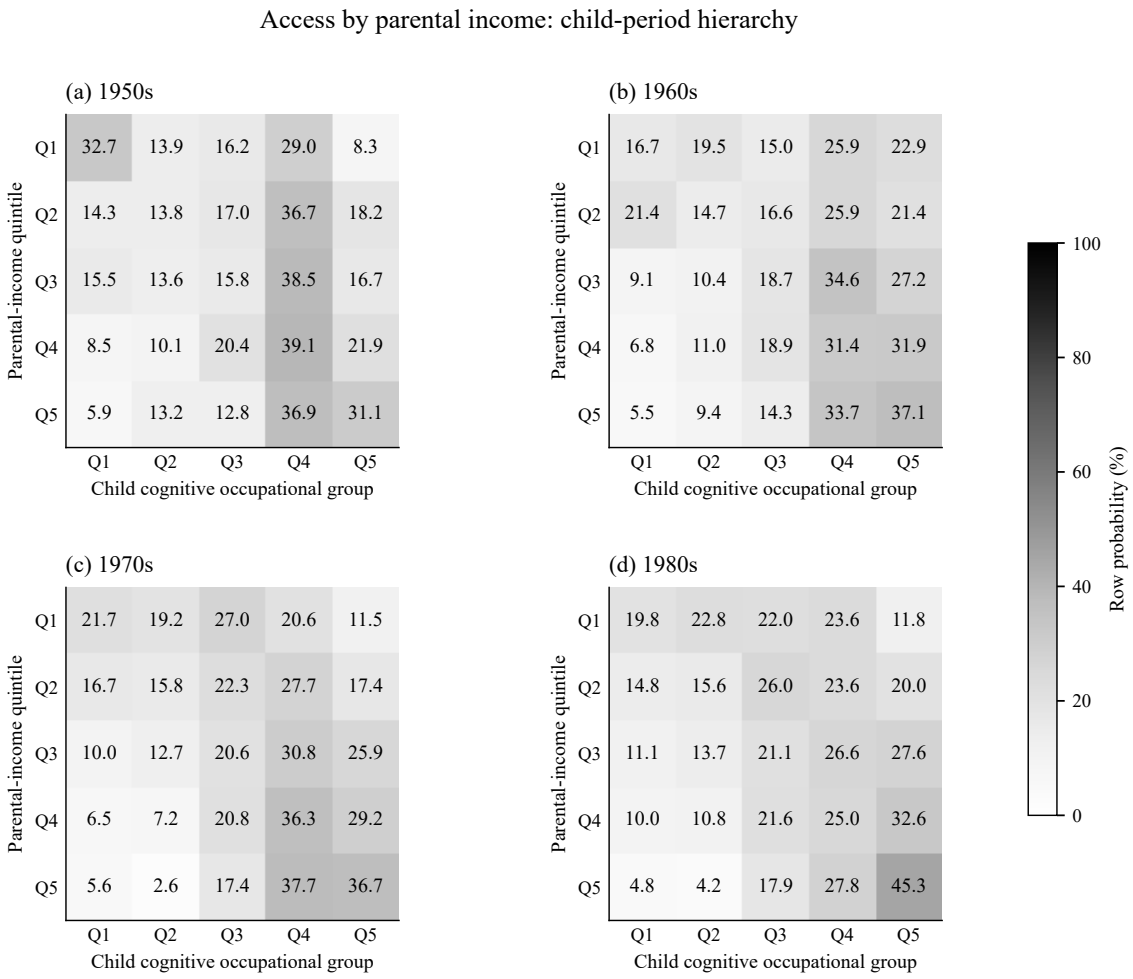


Figure 8: Access to cognitive occupational groups by parental income

Notes: Child-period shared hierarchy, evaluated on the common matrix sample. Rows are parental-income quintiles; columns are child occupational requirement groups. Occupational groups are quintiles of reference occupations, not of workers. Displayed probabilities are descriptive and do not use the IGE regression controls.

11 Conclusion

This paper links occupational skill requirements to intergenerational income persistence across US birth cohorts. Combining linked generations in the PSID with DOT and O*NET descriptors makes it possible to examine occupational requirements over a longer historical span than is available from common individual skill assessments. Cognitive requirements account for the largest positive measured skill contribution in the fixed-profile benchmark. Their contribution depends jointly on the returns associated with cognitively demanding occupations and the relationship between parental income and access to those occupations. Rising returns can therefore coexist with a relatively similar contribution to persistence when the parental-income gradient weakens. This accounting helps explain why changes in the rewards to skills need not translate proportionately into changes in the intergenerational elasticity of income.

The motivating cross-country comparisons reinforce the importance of timing. Adult-period productivity growth is associated with greater persistence, whereas the childhood-aligned comparison does not preserve that positive relationship. The positive association between productivity growth and the PISA parental-background gradient motivates a possible family skill-supply response. Alongside the proposition, these patterns frame the empirical findings as a race between changing rewards and intergenerational allocation, while leaving the causal role of technology and parental investment open.

Historical requirement profiles distinguish baseline persistence, skill reevaluation, within-occupation change, and between-occupation change. The dynamic cognitive analysis yields a small measured within-occupation contribution under both detailed occupations and broader occupational families. This finding concerns the association between parental income and income-valued changes in requirements. Changes shared across family backgrounds can alter the content of work without contributing substantially to the parental-income gradient. Moreover, negative between-occupation components coexist with substantial parental-income gaps in access to high-cognitive destinations. Children from disadvantaged origins can gain relative to their parents' occupational positions while remaining less likely to reach the most demanding occupations.

A central priority for future research is better dynamic measurement of skills and job requirements. Repeated assessments of workers, linked to detailed job histories and family records, would allow researchers to distinguish individual skill development from changes in the requirements of the jobs workers enter. Extending the analysis to additional skill dimensions and demographic groups could then establish whether the cognitive patterns documented here generalize. The present results provide an empirical accounting for that research agenda: understanding intergenerational persistence requires measuring both how work changes and how access to its rewards varies with family background.

References

Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier, 2011.

Melanie Arntz, Cäcilia Lipowski, Guido Neidhöfer, and Ulrich Zierahn-Weilage. Computers as stepping stones? technological change and equality of labor market opportunities. *Journal of Labor Economics*, 43(2):503–543, 2025. doi: 10.1086/727490.

David H. Autor and David Dorn. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597, 2013. doi: 10.1257/aer.103.5.1553. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.5.1553.

David H Autor, Frank Levy, and Richard J Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333, 2003.

Gary S Becker and Nigel Tomes. Human capital and the rise and fall of families. *Journal of labor economics*, 4(3, Part 2):S1–S39, 1986.

Gary S. Becker, Scott Duke Kominers, Kevin M. Murphy, and Jörg L. Spenkuch. A theory of intergenerational mobility. *Journal of Political Economy*, 126(S1):S7–S25, 2018. doi: 10.1086/698759.

Jo Blanden, Paul Gregg, and Lindsey Macmillan. Accounting for intergenerational income persistence: Noncognitive skills, ability and education. *Economic Journal*, 117(519): C43–C60, 2007. doi: 10.1111/j.1468-0297.2007.02034.x.

Miles Corak. Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, 27(3):79–102, 2013.

Flávio Cunha, James J. Heckman, and Susanne M. Schennach. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931, 2010. doi: 10.3982/ECTA6551. URL https://doi.org/10.3982/ECTA6551.

David J. Deming. The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4):1593–1640, 2017. doi: 10.1093/qje/qjx022.

Steven N. Durlauf, Andros Kourtellos, and Chih Ming Tan. The great gatsby curve. *Annual Review of Economics*, 14:571–605, 2022. doi: 10.1146/annurev-economics-082321-122703.

Per-Anders Edin, Peter Fredriksson, Martin Nybom, and Björn Öckert. The rising return to noncognitive skill. *American Economic Journal: Applied Economics*, 14(2):78–100, 2022. doi: 10.1257/app.20190199.

Oded Galor and Daniel Tsiddon. Technological progress, mobility, and economic growth. *American Economic Review*, 87(3):363–382, 1997.

Claudia Goldin and Lawrence F. Katz. *The Race between Education and Technology*. Belknap Press of Harvard University Press, 2008.

Catherine Haeck and Jean-William Laliberté. Careers and intergenerational income mobility. *American Economic Journal: Applied Economics*, 17(1):431–458, 2025. doi: 10.1257/app.20230403.

Steven Haider and Gary Solon. Life-cycle variation in the association between current and lifetime earnings. *American Economic Review*, 96(4):1308–1320, 2006. doi: 10.1257/aer.96.4.1308. URL https://www.aeaweb.org/articles?id=10.1257/aer.96.4.1308.

John Hassler and José V. Rodríguez Mora. Intelligence, social mobility, and growth. *American Economic Review*, 90(4):888–908, 2000. doi: 10.1257/aer.90.4.888.

John Hassler, José V. Rodríguez Mora, and Joseph Zeira. Inequality and mobility. *Journal of Economic Growth*, 12(3):235–259, 2007. doi: 10.1007/s10887-007-9019-x.

Santiago Hermo, Miika Päällysaho, David Seim, and Jesse M. Shapiro. Labor market returns and the evolution of cognitive skills: Theory and evidence. *Quarterly Journal of Economics*, 137(4):2309–2361, 2022. doi: 10.1093/qje/qjac022. URL https://doi.org/10.1093/qje/qjac022.

Lawrence F Katz and Kevin M Murphy. Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78, 1992.

Chul-In Lee and Gary Solon. Trends in intergenerational income mobility. *Review of Economics and Statistics*, 91(4):766–772, 2009. doi: 10.1162/rest.91.4.766.

Erik Lindqvist and Roine Vestman. The labor market returns to cognitive and noncognitive ability: Evidence from the swedish enlistment. *American Economic Journal: Applied Economics*, 3(1):101–128, 2011. doi: 10.1257/app.3.1.101. URL https://www.aeaweb.org/articles?id=10.1257/app.3.1.101.

Martin Nybom and Jan Stuhler. Heterogeneous income profiles and lifecycle bias in intergenerational mobility estimation. *Journal of Human Resources*, 51(1):239–268, 2016. doi: 10.3368/jhr.51.1.239.

Martin Nybom and Jan Stuhler. Interpreting trends in intergenerational mobility. *Journal of Political Economy*, 132(8):2531–2570, 2024. doi: 10.1086/729582.

Gary Solon. Intergenerational income mobility in the united states. *American Economic Review*, 82(3):393–408, 1992. URL https://www.jstor.org/stable/2117312.

Alexandra Spitz-Oener. Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2):235–270, 2006. doi: 10.1086/499972. URL https://www.journals.uchicago.edu/doi/10.1086/499972.

Jan Tinbergen. Substitution of graduate by other labour. *Kyklos: international review for social sciences*, 1974.

A Construction Details and Sample Qualifications

A.1 Income windows and occupation coverage

Child log income remains annual. The selected parent's log income is averaged over the existing observations at child ages 15–17. The cleaned parent window occasionally contains more than three observations because of inconsistent recorded child birth years. Parent ages range broadly in the underlying audit; no new age trim was imposed. Likewise, original source-cohort assignments, household-selection rules, and income cleaning were retained. These choices preserve comparability across the completed analyses but require further data review before final publication.

The starting sample has 14,186 recorded child occupations, of which 14,024 map to a harmonized Census code and 13,940 to civilian occupations. Selected father records number 9,019, with 8,275 recorded occupations and 8,205 matches; mother records number 2,491, with 1,944 recorded occupations and 1,935 matches. These parent counts are child–parent–income-year links, not unique parents. Siblings can repeat parental records. Military and nonoccupation codes do not receive civilian requirement scores, and missing jobs do not receive zero requirements.

The fixed O*NET benchmark retains 13,738 child-years and 4,445 children. The historical joint-skill sample retains 12,985 and 4,161. In that historical sample,1,640 child-years have at least one original income-window record without a matched parent skill profile. Income continues to use all original eligible income records, while parent skill exposure uses the matched subset. Partial windows are concentrated in the 1970s cohort. The later dated extension holds this original matched parent-record set fixed and requires all of its records to be scoreable in each comparison construction. It does not recover previously unmatched parent records or eliminate the inherited difference between income and skill dates.

A.2 Descriptor aggregation and principal components

For O*NET specialty d, SOC s, and descriptor k, first compute the mean of usable specialty ratings $\bar{X}_{sk}$. For common occupation j, the employment-weighted available-constituent descriptor is

$$X_{jk} = \frac{\sum_{s \mapsto j} E_s \bar{X}_{sk} D_{sk}}{\sum_{s \mapsto j} E_s D_{sk}}, \qquad C_{jk} = \frac{\sum_{s \mapsto j} E_s D_{sk}}{\sum_{s \mapsto j} E_s}, \tag{19}$$

where D_{sk} indicates an observed usable descriptor and C_{jk} records employment coverage. A missing constituent descriptor reduces coverage rather than disappearing from its denominator. Coverage refers to the mapped reference inventory, not the entire population of US jobs. OES

employment omits self-employment and does not fully represent agricultural work. No minimum 80 or 95 percent coverage restriction is introduced into the benchmark after inspecting outcomes. Available-constituent means therefore assume that observed profiles represent their common group; some minimum coverage values are low and qualify interpretation.

On the reference occupation set $\mathcal{J}_b$, descriptor standardization and the oriented PC 1 score are

$$Z_{jk}^b = \frac{X_{jk} - \mu_{bk}}{\sigma_{bk}}, \qquad S_j^b = \frac{\sum_k v_{bk} Z_{jk}^b - \bar{H}_b}{s_{H,b}}. \tag{20}$$

Here v_b is the leading correlation-matrix eigenvector, oriented so its coefficient sum is positive; $\bar{H}_b$ and $s_{H,b}$ normalize the resulting PC 1. Reference descriptors use sample SDs. Employment weights operate within occupations, whereas each reference occupation has one observation in PCA. Scores are not weighted by PSID occupational frequencies and are not standardized separately for parents, children, or birth cohorts. Separate domain PCAs do not make the resulting domains mutually orthogonal.

The fixed benchmark uses all 60 descriptors on a common 456-occupation reference. Historical direct PCA has three cognitive DOT columns and seven cognitive O*NET columns: verbal or reasoning inputs are not duplicated in DOT to match the larger O*NET battery. The five historical noncognitive inputs proxy leadership, social perceptiveness/dealing with people, persuasion/influencing people, stress tolerance, and adaptability/variety. The three manual inputs proxy manual dexterity, finger dexterity, and limb coordination. DOT temperaments are presence indicators rather than measurements on O*NET's Importance scale. Source standardization does not establish measurement invariance across these different instruments.

The multiple-date extension reuses the DOT 1977 and O*NET 15.0 cognitive scoring transformations. It does not pool all raw DOT and O*NET descriptors in one PCA or independently normalize every archive release. Known-weight but unrated O*NET cells remain in coverage denominators; a target with unsupported employment weights is withheld. Missing taxonomy lineage and cross-common splits remain unsupported. Legacy and synthetic analyst-transition ratings are retained with flags, not labeled independent contemporaneous observations.

DOT 1977 uses ICPSR 07845 Part 1, whose 60,441 CPS-linked records cover 3,885 observed valid DOT jobs and retain ordinary Census 1970 codes and the original CPS weights. Its expanded-category Part 2 is not merged as if it were the ordinary Census classification. DOT 1991 uses the 12,741-job definition master in ICPSR 06100 Part 3 and the documented old-to-new conversion relation. DOT 1965 uses digitized aptitude records and the published DOT 1965-to-Census 1970 job crosswalk. The historical employment reference is not a measure of employment in every DOT year. O*NET archives, official taxonomy relations, and

Table A.1: Income-year assignment of dated cognitive profiles

Income year	Requirement profile
1967–1976	DOT 1965
1977–1990	DOT 1977
1991–2002	DOT 1991
2003–2004	O*NET 5.1
2005	O*NET 9.0
2006–2009	O*NET 10.0
2010	O*NET 15.0
2011	O*NET 16.0
2012–2014	O*NET 17.0
2015–2018	O*NET 20.0
2019	O*NET 24.1
2020–2021	O*NET 25.0
2022	O*NET 27.1

Notes: The cutoff is December 31 of the income year, not the next reporting survey. Each generation and parent record is dated separately. Some listed years have no PSID record under the biennial schedule. O*NET 8.0 is included in release-history diagnostics but is superseded by 9.0 at the 2005 cutoff. Archive release dates are not universal descriptor rating dates.

BLS occupational employment data supply the later descriptors and aggregation references.[3]

A.3 Profiles, ranks, and descriptive distribution statistics

Rank stability uses Spearman correlations with average ranks and Kendall tau-b with ties. Main comparisons balance detailed occupations within source blocks; pairwise-available rank comparisons and their observation counts are retained separately. The main 22-family profiles preserve the previously estimated family means, allowing the detailed members with valid data to vary. A separate sensitivity restricts detailed membership to the balanced block before aggregating raw descriptors and applying unchanged scores. Even that restriction does not necessarily fix native specialty composition.

The equal-weight distribution describes reference occupations or families. The employment-weight sensitivity fixes anchor occupation totals across all snapshots: original CPS reference totals associated with DOT 1977 and May 2018 OES totals associated with O*NET 15.0. These are not contemporaneous occupational shares or PSID survey weights. Weighted SD is

[3]Source documentation is available from https://www.onetcenter.org/db_releases.html, https://www.onetcenter.org/taxonomy.html, and https://usa.ipums.org/usa-action/variables/OCC2010. Exact local source identities and file hashes accompany the replication materials.

Table A.2: Complete common-occupation cognitive profiles by snapshot

Snapshot	Source block	Complete groups	Minimum score	Maximum score
O*NET 5.1	ONET	422	-4.633	4.006
O*NET 8.0	ONET	417	-3.952	3.446
O*NET 9.0	ONET	417	-3.952	3.207
O*NET 10.0	ONET	419	-3.871	3.207
O*NET 15.0	ONET	440	-2.692	3.637
O*NET 16.0	ONET	443	-2.692	3.633
O*NET 17.0	ONET	444	-2.692	3.633
O*NET 20.0	ONET	446	-2.692	3.633
O*NET 24.1	ONET	449	-2.648	3.633
O*NET 25.0	ONET	449	-2.648	3.430
O*NET 27.1	ONET	435	-2.648	3.430
DOT 1965	DOT	267	-1.729	2.162
DOT 1977	DOT	267	-1.871	2.502
DOT 1991	DOT	267	-1.889	2.523

Notes: Counts refer to common occupation profiles with all required cognitive descriptors, not to raw source job titles or PSID observations. Source support can vary over time. No global balanced universe is imposed on the main dated linkage. The profile rank and trend diagnostic separately uses balanced support.

a descriptive population moment, $[\sum_j \omega_j (S_j - \bar{S})^2]^{1/2}$ with normalized weights. The weighted median is the smallest observed value reaching cumulative weight one half. The median of changes need not equal the change in medians. Increasing or unchanged scores use a numerical tolerance of 10^{-10}. No conventional iid occupation-level significance test is reported for these archived reference populations.

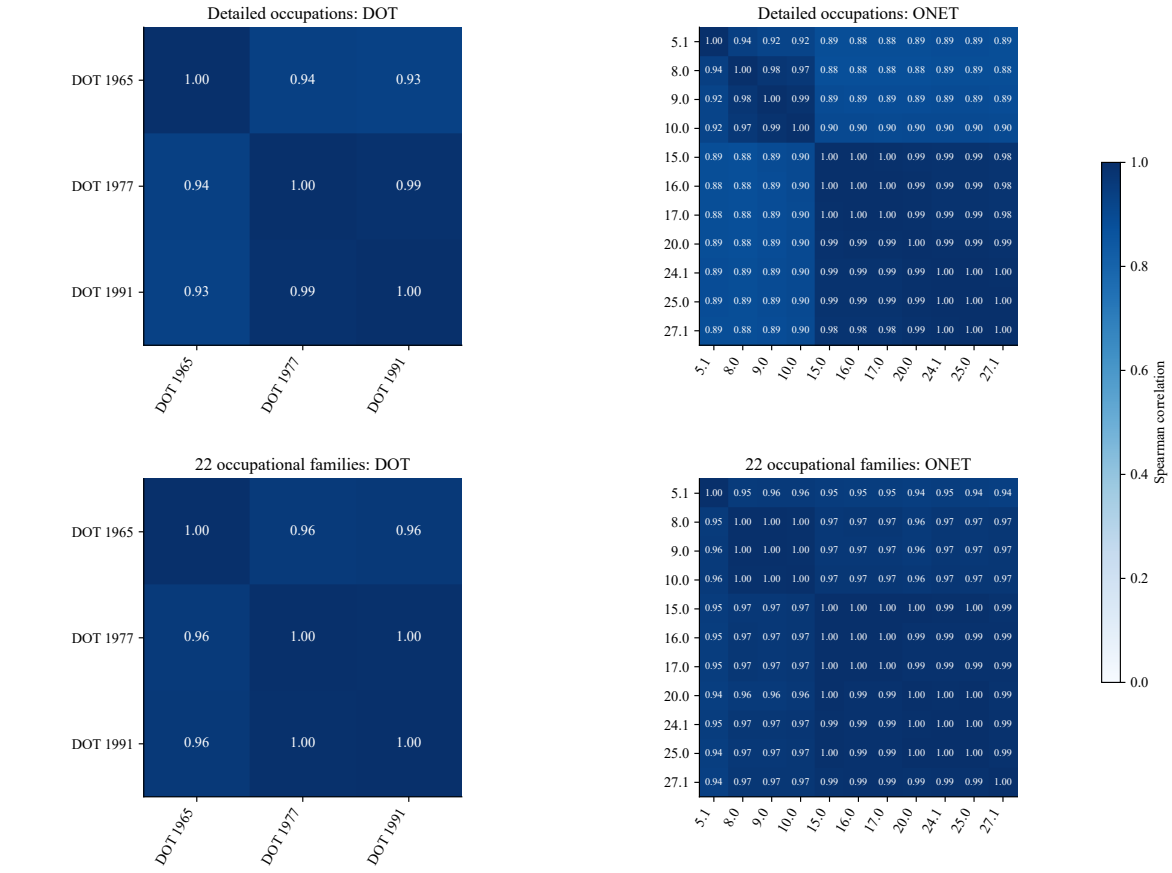


Figure A.1: Stability of occupational cognitive rankings

Notes: Spearman correlations on balanced support within source blocks, using the revised blue palette. Rounded entries of 1.00 need not represent exact equality. These are archive-profile descriptors, without iid occupation significance tests.

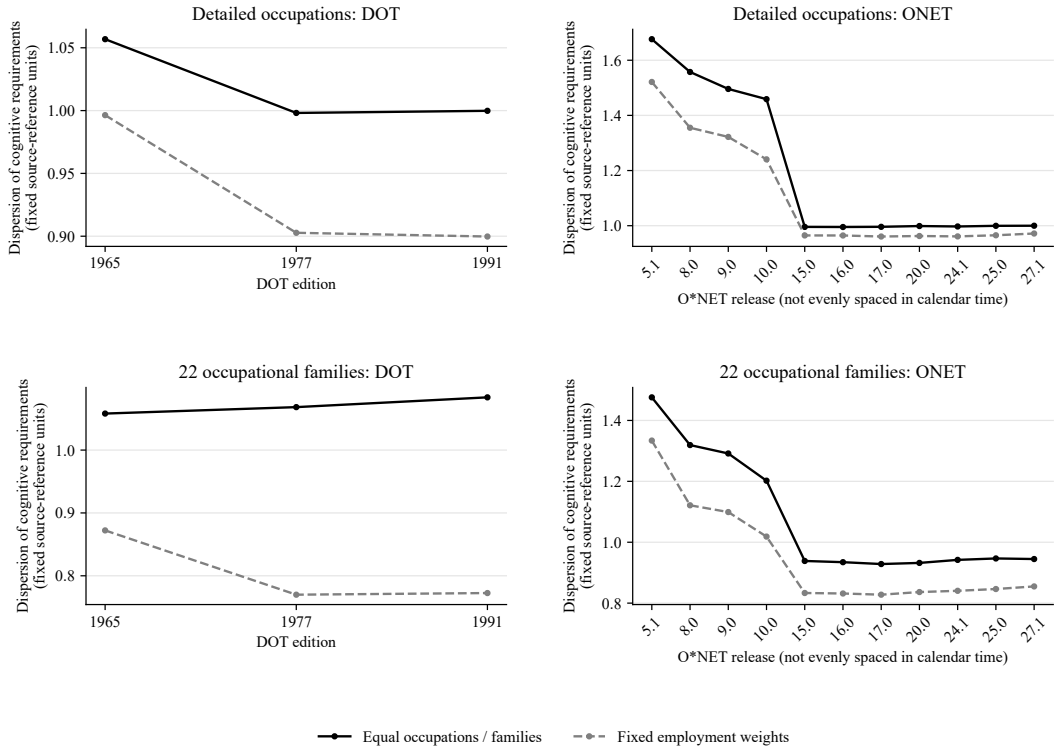
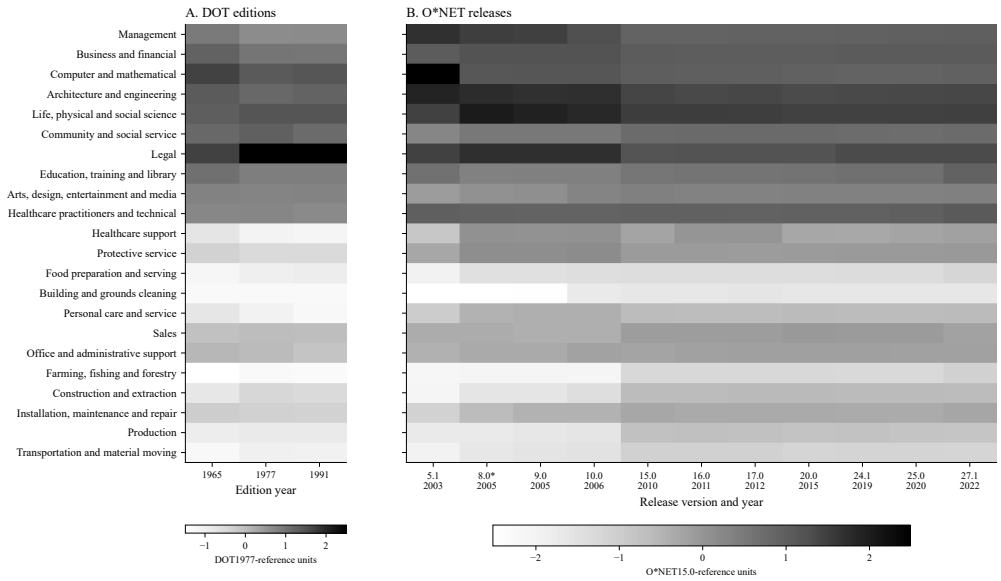


Figure A.2: Dispersion of cognitive requirements within source blocks

Notes: SDs are descriptive population moments on balanced support. Equal occupation and fixed employment weights are shown. No per-wave score normalization. Source blocks and panels retain separate units/scales.



Notes: Rows are the same 22 occupational families in both panels. Darker cells indicate higher cognitive requirements within each source block.
Panels use separate color scales. DOT and O*NET are not cardinally calibrated; cross-panel shades are not directly comparable.
* O*NET8.0 is shown for historical completeness; the approved year-end timing rule selects 9.0 for 2005. Release year is not a universal rating date.

Figure A.3: Cognitive requirements across occupational families and snapshots

Notes: All 22 families and 14 constructed snapshots. Darker cells indicate higher requirements within a source block. The DOT and O*NET panels have separate color scales and are not cardinally calibrated.

A.4 Item-level skill definitions and rating scales

Tables A.3–A.5 list all 60 inputs in the broad fixed O*NET benchmark, with the archived element identifier, source domain, and scale. These are the descriptors actually used to form the indices, not illustrative examples. Table A.6 reports the narrower common-core historical mapping. The distinction matters: the same Social Perceptiveness descriptor uses Level in the broad fixed benchmark but Importance in the historical interpersonal battery. No Importance-by-Level product enters either construction.

O*NET Level (LV) indicates the amount or complexity of a requirement and ranges from 0 to 7, with descriptor-specific anchors. Importance (IM) indicates its occupational importance and ranges from 1 (not important) to 5 (extremely important). We use the archived database ratings rather than the 0–100 presentation scores displayed on the O*NET website.[4] Neither numerical range is a directly observed quantity of an individual's human capital, and identical nominal ranges do not establish invariant anchors or rating methods across releases.

DOT aptitude categories have the opposite raw direction: category 1 denotes a higher requirement than category 5. We orient them using $6-X$, so higher transformed values denote higher requirements. General Learning is observed on categories 1–4 in the audited DOT files; we do not manufacture a fifth observed category. DOT temperaments are coded as the presence or absence of the relevant work requirement, with a wholly unknown temperament record left missing. The historical cross-source correspondences are conceptual approximations: general learning is not separately observed deductive and inductive reasoning, nor does one verbal aptitude rating supply four independent verbal inputs. Thus DOT cognition enters PCA with three unique inputs, and O*NET cognition with seven.

The manual core pairs DOT manual dexterity, finger dexterity, and eye–hand–foot coordination with O*NET Manual Dexterity, Finger Dexterity, and Multilimb Coordination. This core measures dexterity and coordination, whereas the broad fixed manual index also includes strength, stamina, flexibility, and equilibrium. The historical noncognitive core combines leadership/direction, dealing with people, influencing people, stress, and variety/change requirements. It is a broad occupational interpersonal and self-regulatory construct, not a test score for personality or an estimate of a Big Five trait.

Official suppression flags take priority and leave ratings missing. The implemented benchmark treats a nonsuppressed Level observation marked not relevant as zero, an explicit coding assumption rather than a claim that the published numeric rating was zero. Other absent observations remain missing. A prepared measurement sensitivity retains the published non-

[4] See https://www.onetonline.org/help/online/scales. The archived releases, rather than subsequent revisions of the O*NET content model, define the measures used here.

suppressed not-relevant ratings; its availability is not presented as a completed decomposition test. There is no available-item substitution across descriptors in the primary PCA reference.

We orient each component toward higher average requirements without forcing every loading to be positive. For example, Spatial Orientation has a negative loading in the broad cognitive PC1 and is not individually reversed or deleted. Table A.8 reports the source-specific fit of the benchmark PC1s. The relatively low historical DOT noncognitive explained variance is a measurement qualification; higher noncognitive PCs were explored in preparation, but are not inserted into the reported baseline decomposition. Table A.7 identifies the exact memberships of the narrower fixed-profile and five-domain alternatives that were estimated.

B Additional Income Decompositions

The following exhibits retain the distinction between alternative measurement and alternative estimands. The narrow fixed-profile model uses 14 cognitive, seven self-regulation, and 19 manual inputs. The five-domain model replaces the broad noncognitive and manual indices with character, social, psychomotor, and physical PC 1 indices; it does not include an aggregate alongside its own subdomains. These models retain the same 13,738 child-years as the broad fixed-profile benchmark. Their combined shares differ little, but social and character associations, and psychomotor and physical associations, can offset. Appendix Table B.1 reports combined contributions rather than imposing that a sum of subdomain products equal a broad-domain product.

The three-snapshot joint-skill decomposition remains a separate model from the dated cognitive-only extension. Table B.2 reports its aggregate accounting. Table B.3 reports cognitive components for father-selected pairs, DOT 1991 historical replacement, and the parental-occupation counterfactual. The counterfactual is compared against the baseline reestimated on the same additional-profile sample. Restricting to fathers keeps eligible daughters as well as sons; no unreported male-only sample is implied.

B.1 Standardized log incomes

Standardization uses each generation's own cohort-specific log-income mean and sample SD on the exact regression sample, with the same repeated-child-year weighting. The bootstrap reestimates these moments. This produces conditional standardized slopes, not unconditional Pearson correlations. For a fixed child-skill projection the component shares are invariant to income units. In the generation-specific accounting, the parent-return normalization leaves baseline persistence invariant and reallocates the sum of baseline and reevaluation according

to equation (17). Appendix Table B.4 reports the dated cognitive model rather than treating these components as if all had the same scale factor.

B.2 Additional numerical results

Table B.6 reports both generations' cognitive income coefficients, their difference, and the child's parental-income gradient in the dated single-skill specification. Table B.7 retains the net cognitive, noncognitive, and manual contributions from the joint three-snapshot total-income model. This makes the empirical basis for emphasizing cognition visible without equating the smaller domains to zero. Table B.8 shows the separate fixed-profile character, social, psychomotor, and physical products that can be concealed by the broad indices; Figures B.4 and B.5 display their coefficients and gradients.

The 22-family dated exercise changes predictors and reestimates returns, so it can change baseline, reevaluation, and the remainder as well as within and between. Table B.9 retains this full accounting. Tables B.10 and B.11 report the fixed-employment-weight and balanced-detailed-constituent profile diagnostics, respectively. Figure B.6 complements the full family-by-wave heatmap with the earlier profile-line format. These exhibits use existing saved estimates and do not introduce additional regressions or a new occupational requirement scale.

Table A.3: Fixed O*NET benchmark: cognitive inputs

Element ID	Descriptor	Domain block	Scale
1.A.1.a.1	Oral Comprehension	Cognitive abilities	LV: 0–7
1.A.1.a.2	Written Comprehension	Cognitive abilities	LV: 0–7
1.A.1.a.3	Oral Expression	Cognitive abilities	LV: 0–7
1.A.1.a.4	Written Expression	Cognitive abilities	LV: 0–7
1.A.1.b.1	Fluency of Ideas	Cognitive abilities	LV: 0–7
1.A.1.b.2	Originality	Cognitive abilities	LV: 0–7
1.A.1.b.3	Problem Sensitivity	Cognitive abilities	LV: 0–7
1.A.1.b.4	Deductive Reasoning	Cognitive abilities	LV: 0–7
1.A.1.b.5	Inductive Reasoning	Cognitive abilities	LV: 0–7
1.A.1.b.6	Information Ordering	Cognitive abilities	LV: 0–7
1.A.1.b.7	Category Flexibility	Cognitive abilities	LV: 0–7
1.A.1.c.1	Mathematical Reasoning	Cognitive abilities	LV: 0–7
1.A.1.c.2	Number Facility	Cognitive abilities	LV: 0–7
1.A.1.d.1	Memorization	Cognitive abilities	LV: 0–7
1.A.1.e.1	Speed of Closure	Cognitive abilities	LV: 0–7
1.A.1.e.2	Flexibility of Closure	Cognitive abilities	LV: 0–7
1.A.1.e.3	Perceptual Speed	Cognitive abilities	LV: 0–7
1.A.1.f.1	Spatial Orientation	Cognitive abilities	LV: 0–7
1.A.1.f.2	Visualization	Cognitive abilities	LV: 0–7
1.A.1.g.1	Selective Attention	Cognitive abilities	LV: 0–7
1.A.1.g.2	Time Sharing	Cognitive abilities	LV: 0–7

Notes: Exact archived O*NET 24.1 inputs. LV denotes Level; IM denotes Importance. All inputs are aggregated to common occupations before reference standardization and PCA. This inventory has 21 unique descriptors. Higher raw O*NET ratings denote greater requirements; orientation of the aggregate PC1 does not force each individual loading positive.

Table A.4: Fixed O*NET benchmark: broad noncognitive inputs

Element ID	Descriptor	Domain block	Scale
2.B.1.a	Social Perceptiveness	Social Skills	LV: 0–7
2.B.1.b	Coordination	Social Skills	LV: 0–7
2.B.1.c	Persuasion	Social Skills	LV: 0–7
2.B.1.d	Negotiation	Social Skills	LV: 0–7
2.B.1.e	Instructing	Social Skills	LV: 0–7
2.B.1.f	Service Orientation	Social Skills	LV: 0–7
1.C.1.a	Achievement/Effort	Work Styles	IM: 1–5
1.C.1.b	Persistence	Work Styles	IM: 1–5
1.C.1.c	Initiative	Work Styles	IM: 1–5
1.C.2.b	Leadership	Work Styles	IM: 1–5
1.C.3.a	Cooperation	Work Styles	IM: 1–5
1.C.3.b	Concern for Others	Work Styles	IM: 1–5
1.C.3.c	Social Orientation	Work Styles	IM: 1–5
1.C.4.a	Self Control	Work Styles	IM: 1–5
1.C.4.b	Stress Tolerance	Work Styles	IM: 1–5
1.C.4.c	Adaptability/Flexibility	Work Styles	IM: 1–5
1.C.5.a	Dependability	Work Styles	IM: 1–5
1.C.5.b	Attention to Detail	Work Styles	IM: 1–5
1.C.5.c	Integrity	Work Styles	IM: 1–5
1.C.6	Independence	Work Styles	IM: 1–5

Notes: Exact archived O*NET 24.1 inputs. LV denotes Level; IM denotes Importance. All inputs are aggregated to common occupations before reference standardization and PCA. This inventory has 20 unique descriptors. Higher raw O*NET ratings denote greater requirements; orientation of the aggregate PC1 does not force each individual loading positive.

Table A.5: Fixed O*NET benchmark: manual inputs

Element ID	Descriptor	Domain block	Scale
1.A.2.a.1	Arm-Hand Steadiness	Psychomotor	LV: 0–7
1.A.2.a.2	Manual Dexterity	Psychomotor	LV: 0–7
1.A.2.a.3	Finger Dexterity	Psychomotor	LV: 0–7
1.A.2.b.1	Control Precision	Psychomotor	LV: 0–7
1.A.2.b.2	Multilimb Coordination	Psychomotor	LV: 0–7
1.A.2.b.3	Response Orientation	Psychomotor	LV: 0–7
1.A.2.b.4	Rate Control	Psychomotor	LV: 0–7
1.A.2.c.1	Reaction Time	Psychomotor	LV: 0–7
1.A.2.c.2	Wrist-Finger Speed	Psychomotor	LV: 0–7
1.A.2.c.3	Speed of Limb Movement	Psychomotor	LV: 0–7
1.A.3.a.1	Static Strength	Physical	LV: 0–7
1.A.3.a.2	Explosive Strength	Physical	LV: 0–7
1.A.3.a.3	Dynamic Strength	Physical	LV: 0–7
1.A.3.a.4	Trunk Strength	Physical	LV: 0–7
1.A.3.b.1	Stamina	Physical	LV: 0–7
1.A.3.c.1	Extent Flexibility	Physical	LV: 0–7
1.A.3.c.2	Dynamic Flexibility	Physical	LV: 0–7
1.A.3.c.3	Gross Body Coordination	Physical	LV: 0–7
1.A.3.c.4	Gross Body Equilibrium	Physical	LV: 0–7

Notes: Exact archived O*NET 24.1 inputs. LV denotes Level; IM denotes Importance. All inputs are aggregated to common occupations before reference standardization and PCA. This inventory has 19 unique descriptors. Higher raw O*NET ratings denote greater requirements; orientation of the aggregate PC1 does not force each individual loading positive.

Table A.6: Historical skill batteries and conceptual source correspondences

Domain	O*NET descriptor	Element ID	O*NET scale	DOT counterpart
Cognitive	Deductive Reasoning	1.A.1.b.4	LV: 0–7	General learning / intelligence
Cognitive	Inductive Reasoning	1.A.1.b.5	LV: 0–7	Same general-learning input
Cognitive	Oral Comprehension	1.A.1.a.1	LV: 0–7	Verbal aptitude
Cognitive	Written Comprehension	1.A.1.a.2	LV: 0–7	Same verbal input
Cognitive	Oral Expression	1.A.1.a.3	LV: 0–7	Same verbal input
Cognitive	Written Expression	1.A.1.a.4	LV: 0–7	Same verbal input
Cognitive	Number Facility	1.A.1.c.2	LV: 0–7	Numerical aptitude
Noncognitive	Leadership	1.C.2.b	IM: 1–5	D: directing / controlling / planning
Noncognitive	Social Perceptiveness	2.B.1.a	IM: 1–5	P: dealing with people
Noncognitive	Persuasion	2.B.1.c	IM: 1–5	I: influencing people
Noncognitive	Stress Tolerance	1.C.4.b	IM: 1–5	S: stress requirements
Noncognitive	Adaptability/Flexibility	1.C.4.c	IM: 1–5	V: variety and change
Manual	Manual Dexterity	1.A.2.a.2	LV: 0–7	Manual dexterity
Manual	Finger Dexterity	1.A.2.a.3	LV: 0–7	Finger dexterity
Manual	Multilimb Coordination	1.A.2.b.2	LV: 0–7	Eye–hand–foot coordination

Notes: Three unique DOT aptitudes correspond to seven distinct O*NET cognitive inputs; repeated descriptions in this table do not create duplicate DOT PCA columns. DOT aptitude raw categories run from 1 (higher requirement) to 5 (lower), oriented using six minus raw; General Learning is observed on 1–4. The five DOT temperament inputs are 0/1 presence indicators. This is a conceptual bridge, not a cardinal calibration. Historical social descriptors use Importance, whereas the broad fixed social battery uses Level. DOT 1965 enters the extended cognitive-only analysis, not a three-domain 1965 model.

Table A.7: Exact memberships of estimated fixed-profile alternatives

Index	Input membership	Items
Narrow cognitive	First 14 cognitive items in Table A.3	14
Self-regulation	Persistence; Self Control; Stress Tolerance; Adaptability/Flexibility; Dependability; Attention to Detail; Integrity	7
Character / work styles	All Work Styles in Table A.4	14
Social	All Social Skills in Table A.4	6
Psychomotor	All 1.A.2 inputs in Table A.5	10
Physical	All 1.A.3 inputs in Table A.5	9

Notes: Narrow cognition includes Oral Comprehension through Memorization in the displayed cognitive inventory. Self-regulation uses the named seven Work Styles inputs. The narrow model includes narrow cognitive, self-regulation, and the broad 19-input manual index. The five-domain model includes broad cognitive, character, social, psychomotor, and physical scores jointly. PC2 combinations and balanced-facet scores were not substituted into the reported decompositions.

Table A.8: PC1 explained variance in historical benchmark batteries

Source	Reference occupations	Cognitive (%)	Noncognitive (%)	Manual (%)
DOT 1977	267	89.8	34.4	58.1
O*NET 15.0	423	86.5	69.2	82.2
O*NET 24.1	448	86.4	69.4	87.9
DOT 1991	267	89.9	36.7	57.9

Notes: Direct descriptor PCA. The broad fixed 60-item benchmark separately explains 62.1, 54.3, and 75.1 percent in its 456-occupation reference. These are different batteries, so higher historical cognitive variance shares do not establish higher validity. The dated extension applies frozen DOT 1977 and O*NET 15.0 coefficients instead of estimating a separate PCA for every release.

Table B.1: Alternative fixed O*NET skill definitions

Cohort	Definition	IGE	All skills	Remainder	Skill share (%)
1950s	Broad	0.363	0.134	0.229	36.8
1950s	Narrow	0.363	0.136	0.227	37.5
1950s	Five domains	0.363	0.137	0.226	37.7
1960s	Broad	0.374	0.098	0.277	26.1
1960s	Narrow	0.374	0.099	0.275	26.5
1960s	Five domains	0.374	0.100	0.274	26.7
1970s	Broad	0.391	0.150	0.241	38.4
1970s	Narrow	0.391	0.154	0.237	39.4
1970s	Five domains	0.391	0.153	0.238	39.1
1980s	Broad	0.347	0.121	0.226	34.9
1980s	Narrow	0.347	0.123	0.224	35.4
1980s	Five domains	0.347	0.119	0.227	34.5
1990s†	Broad	0.506	0.123	0.383	24.4
1990s†	Narrow	0.506	0.125	0.381	24.7
1990s†	Five domains	0.506	0.127	0.379	25.1

Notes: The broad, narrow and five-domain models use the same 13,738 child-years. Five domains are cognitive, character, social, psychomotor and physical. Individual allocations can differ even when the combined product sum is similar. Shares are point-estimate accounting shares rather than causal mediation percentages.

Table B.2: Three historical snapshots: joint three-domain total-income decomposition

Cohort	IGE	Baseline	Reevaluation	Within	Between	Remainder
1950s	0.365 [0.287, 0.439]	0.195 [0.149, 0.243]	-0.003 [-0.043, 0.034]	0.0000 [0.0000, 0.0000]	-0.065 [-0.099, -0.032]	0.238 [0.168, 0.305]
1960s	0.364 [0.294, 0.435]	0.245 [0.207, 0.290]	-0.038 [-0.079, -0.004]	-0.0003 [-0.0008, -0.0000]	-0.102 [-0.130, -0.075]	0.260 [0.196, 0.328]
1970s	0.431 [0.352, 0.512]	0.173 [0.129, 0.231]	0.035 [-0.002, 0.067]	-0.0010 [-0.0038, 0.0014]	-0.057 [-0.102, -0.013]	0.281 [0.209, 0.351]
1980s	0.338 [0.275, 0.404]	0.140 [0.102, 0.184]	0.015 [-0.015, 0.044]	-0.0019 [-0.0058, 0.0003]	-0.030 [-0.070, 0.010]	0.215 [0.159, 0.272]
1990s†	0.618 [0.190, 0.999]	0.186 [0.094, 0.316]	0.074 [-0.063, 0.192]	0.0017 [-0.0011, 0.0066]	-0.159 [-0.307, -0.002]	0.515 [0.098, 0.909]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Components are separately projected on parental income with the same IGE controls; the remainder is obtained by subtraction. Rounding can affect displayed add-up.

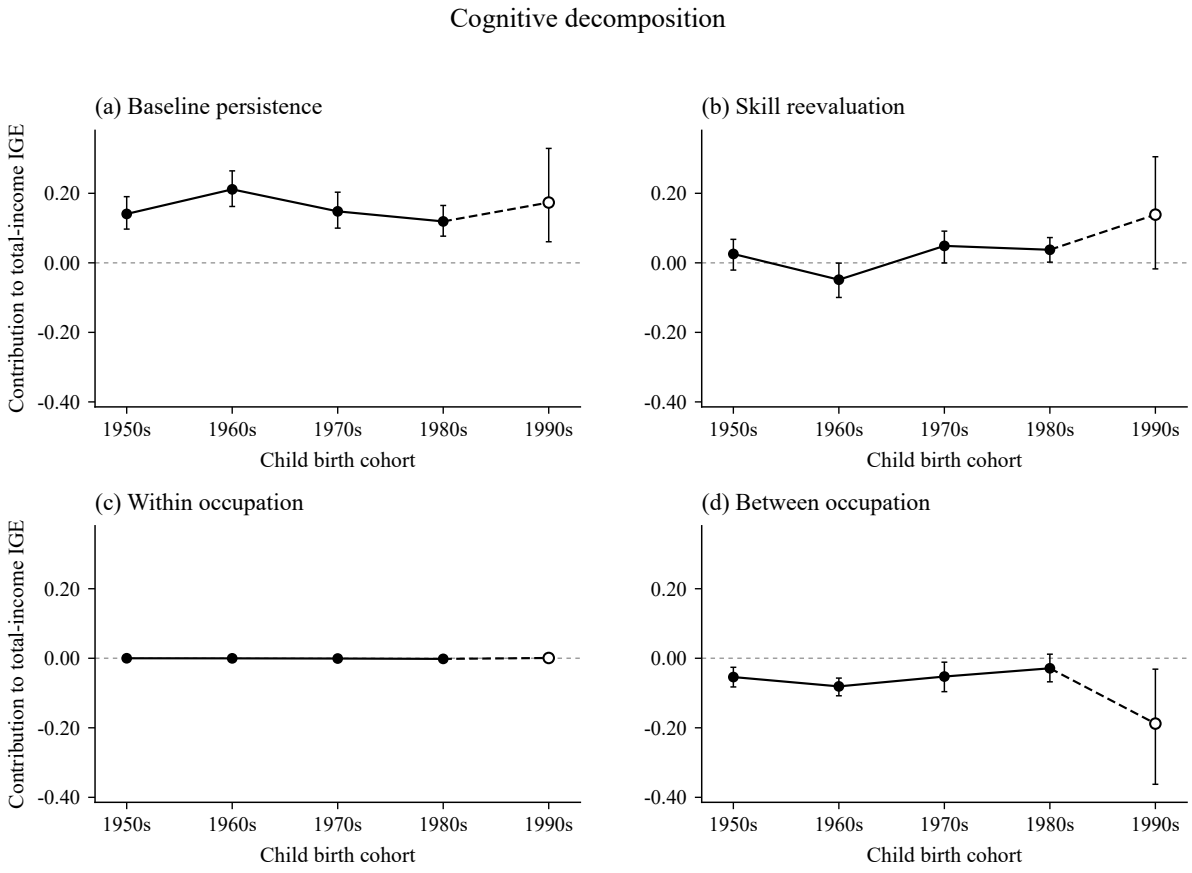


Figure B.1: Cognitive components in the three-domain historical model

Notes: Cognitive terms from the joint cognitive/noncognitive/manual model on 12,985 child-years. These are distinct from the dated single-skill estimates. Intervals condition on source profiles.

Table B.3: Historical cognitive within and between sensitivity checks

Specification	Cohort	Within cognition	Between cognition
Father selected	1950s	-0.0000 [-0.0000, 0.0000]	-0.0395 [-0.0687, -0.0113]
Father selected	1960s	-0.0004 [-0.0009, -0.0000]	-0.0672 [-0.0952, -0.0426]
Father selected	1970s	-0.0010 [-0.0039, 0.0014]	-0.0784 [-0.1265, -0.0292]
Father selected	1980s	-0.0025 [-0.0082, 0.0005]	-0.0479 [-0.0971, 0.0020]
Father selected	1990s†	0.0002 [-0.0025, 0.0027]	-0.1623 [-0.4139, 0.0999]
DOT 1991	1950s	-0.0000 [-0.0000, 0.0000]	-0.0538 [-0.0809, -0.0259]
DOT 1991	1960s	-0.0003 [-0.0007, -0.0000]	-0.0878 [-0.1151, -0.0616]
DOT 1991	1970s	-0.0011 [-0.0037, 0.0012]	-0.0567 [-0.1004, -0.0182]
DOT 1991	1980s	-0.0017 [-0.0049, 0.0005]	-0.0293 [-0.0728, 0.0094]
DOT 1991	1990s†	0.0008 [-0.0017, 0.0039]	-0.1879 [-0.3630, -0.0300]
Common counterfactual sample	1950s	-0.0000 [-0.0000, 0.0000]	-0.0541 [-0.0846, -0.0247]
Common counterfactual sample	1960s	-0.0003 [-0.0006, -0.0000]	-0.0801 [-0.1058, -0.0572]
Common counterfactual sample	1970s	-0.0007 [-0.0031, 0.0014]	-0.0557 [-0.0962, -0.0185]
Common counterfactual sample	1980s	-0.0018 [-0.0055, 0.0005]	-0.0247 [-0.0668, 0.0123]
Common counterfactual sample	1990s†	0.0009 [-0.0015, 0.0037]	-0.1921 [-0.3615, -0.0253]
Parental-occupation counterfactual	1950s	-0.0000 [-0.0000, 0.0000]	-0.0541 [-0.0817, -0.0254]
Parental-occupation counterfactual	1960s	0.0014 [-0.0011, 0.0038]	-0.0818 [-0.1076, -0.0576]
Parental-occupation counterfactual	1970s	-0.0026 [-0.0175, 0.0118]	-0.0539 [-0.0905, -0.0166]
Parental-occupation counterfactual	1980s	-0.0014 [-0.0111, 0.0086]	-0.0252 [-0.0602, 0.0130]
Parental-occupation counterfactual	1990s†	0.0105 [-0.0153, 0.0380]	-0.2017 [-0.3847, -0.0413]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Cognitive terms come from the joint three-domain model. Father selected includes sons and daughters. DOT 1991 replaces historical ratings. The last two specifications use the same additional-profile sample. These are not dated cognitive-only coefficients.

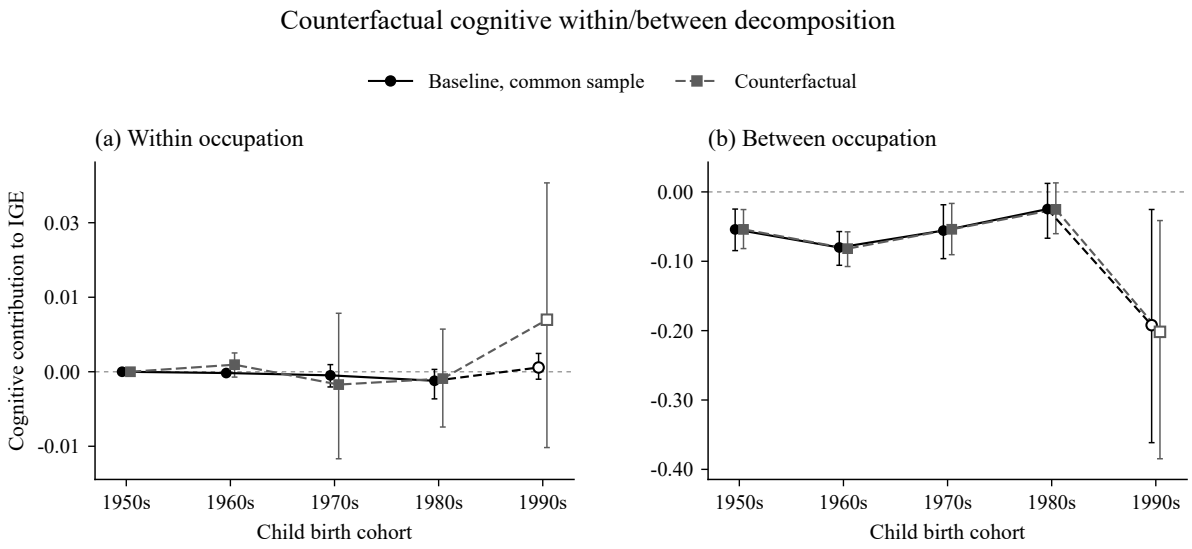


Figure B.2: Parental-occupation counterfactual and realized occupational change

Notes: Baseline and counterfactual compared on identical additional-profile observations. Counterfactual within change evaluates the parental occupation at the child-period profile for stayers and switchers. Differences concern accounting allocations, not identified behavioral responses.

Table B.4: Dated cognition with standardized log incomes

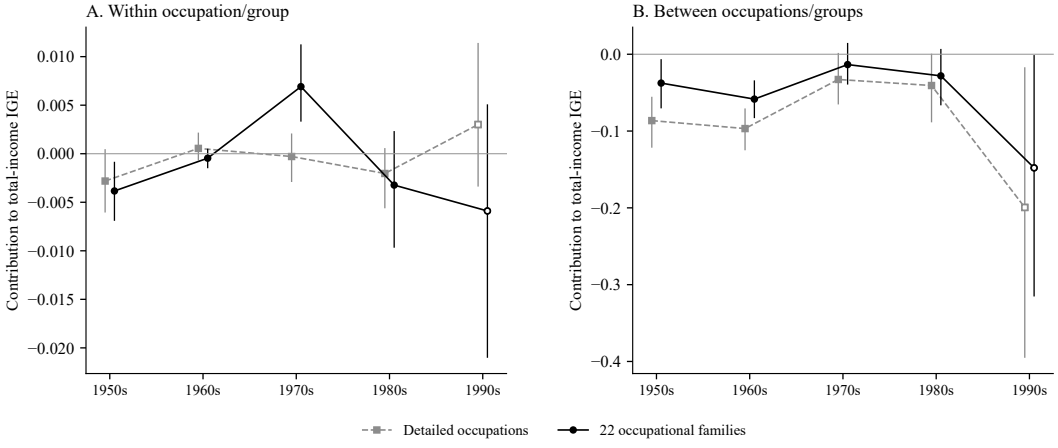
Cohort	IGE	Baseline	Reevaluation	Within	Between	Remainder
1950s	0.308 [0.246, 0.368]	0.197 [0.157, 0.241]	-0.015 [-0.050, 0.013]	-0.0024 [-0.0050, 0.0004]	-0.073 [-0.100, -0.047]	0.202 [0.142, 0.258]
1960s	0.305 [0.254, 0.359]	0.240 [0.196, 0.288]	-0.083 [-0.117, -0.049]	0.0005 [-0.0006, 0.0018]	-0.080 [-0.104, -0.060]	0.228 [0.178, 0.282]
1970s	0.367 [0.311, 0.422]	0.173 [0.126, 0.224]	-0.045 [-0.080, -0.013]	-0.0002 [-0.0025, 0.0018]	-0.028 [-0.055, 0.001]	0.268 [0.216, 0.314]
1980s	0.298 [0.241, 0.350]	0.117 [0.079, 0.158]	0.035 [0.013, 0.053]	-0.0019 [-0.0051, 0.0006]	-0.038 [-0.081, 0.001]	0.185 [0.135, 0.233]
1990s†	0.501 [0.138, 0.807]	0.142 [0.050, 0.263]	0.104 [0.001, 0.225]	0.0025 [-0.0028, 0.0091]	-0.163 [-0.318, -0.014]	0.415 [0.052, 0.727]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Components are separately projected on parental income with the same IGE controls; the remainder is obtained by subtraction. Rounding can affect displayed add-up.

Table B.5: Dated cognitive contributions under occupational aggregation

Cohort	Occupation unit	Within	Between
1950s	Detailed	-0.0028 [-0.0061, 0.0005]	-0.0864 [-0.1217, -0.0554]
1950s	22 families	-0.0038 [-0.0069, -0.0008]	-0.0376 [-0.0705, -0.0064]
1960s	Detailed	0.0005 [-0.0007, 0.0022]	-0.0968 [-0.1250, -0.0706]
1960s	22 families	-0.0005 [-0.0015, 0.0005]	-0.0584 [-0.0832, -0.0340]
1970s	Detailed	-0.0003 [-0.0029, 0.0021]	-0.0328 [-0.0654, 0.0017]
1970s	22 families	0.0069 [0.0033, 0.0113]	-0.0135 [-0.0395, 0.0147]
1980s	Detailed	-0.0020 [-0.0056, 0.0006]	-0.0407 [-0.0887, 0.0007]
1980s	22 families	-0.0032 [-0.0097, 0.0023]	-0.0281 [-0.0665, 0.0069]
1990s†	Detailed	0.0030 [-0.0034, 0.0114]	-0.1994 [-0.3952, -0.0169]
1990s†	22 families	-0.0059 [-0.0210, 0.0051]	-0.1479 [-0.3154, -0.0011]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Identical 12,481 child-years and 4,024 children. Family profiles aggregate raw descriptors before unchanged scoring, and returns are reestimated. Within means equality at the stated occupational level. The 1950s family within interval is negative; the 1970s interval is positive. Thus small magnitudes do not imply universal statistical insignificance.



Notes: Dated cognitive requirements in both specifications; same 12,481 child-years and original parental weights.
Bars: 95% paired child-cluster bootstrap intervals, 999 draws. Hollow markers: 1990s age-33 extrapolation.
Family raw requirements are aggregated before frozen PCA scoring; prices are reestimated. Panels have separate vertical scales.

Figure B.3: Dated cognitive contributions: detailed occupations and 22 families

Notes: Identical sample, income definitions, date assignments and parent-record weights. Family-level raw profiles are constructed before scoring; returns are reestimated. Bars use paired child-cluster bootstrap draws.

Table B.6: Generation-specific cognitive coefficients and intergenerational requirement gradients

Cohort	Parent τ_c^p	Child τ_c^o	$\Delta\tau_c$	Child gradient π_c
1950s	0.224 [0.197, 0.253]	0.245 [0.212, 0.279]	0.021 [-0.022, 0.064]	0.515 [0.417, 0.622]
1960s	0.289 [0.258, 0.322]	0.227 [0.192, 0.265]	-0.062 [-0.104, -0.015]	0.406 [0.312, 0.497]
1970s	0.299 [0.253, 0.348]	0.259 [0.225, 0.297]	-0.040 [-0.100, 0.018]	0.449 [0.355, 0.558]
1980s	0.264 [0.218, 0.310]	0.371 [0.324, 0.415]	0.108 [0.043, 0.171]	0.330 [0.244, 0.411]
1990s†	0.169 [0.099, 0.238]	0.357 [0.233, 0.466]	0.188 [0.043, 0.323]	0.294 [-0.042, 0.678]

Notes: Brackets report 95 percent whole-child bootstrap percentile intervals from 999 draws, conditional on constructed requirements. †: the 1990s age-33 target extrapolates from observed ages 30–32. One cognitive regressor enters each generation-specific return equation. Coefficients are conditional income associations per fixed source-reference score unit. The child requirement gradient is a standardized-level-on-log-parent-income slope, not a literal log-log elasticity.

Table B.7: Net skill contributions in the joint three-snapshot total-income model

Cohort	Cognitive	Noncognitive	Manual
1950s	0.112 [0.083, 0.146]	0.016 [0.001, 0.036]	-0.001 [-0.009, 0.007]
1960s	0.082 [0.055, 0.111]	0.029 [0.013, 0.048]	-0.006 [-0.013, -0.001]
1970s	0.144 [0.109, 0.185]	0.011 [-0.009, 0.032]	-0.004 [-0.018, 0.009]
1980s	0.126 [0.092, 0.168]	0.001 [-0.010, 0.012]	-0.004 [-0.013, 0.005]
1990s†	0.125 [-0.020, 0.314]	-0.019 [-0.078, 0.024]	-0.003 [-0.036, 0.029]

Notes: Brackets report 95 percent whole-child bootstrap percentile intervals from 999 draws, conditional on constructed requirements. †: the 1990s age-33 target extrapolates from observed ages 30–32. Each net term sums its domain-specific baseline, reevaluation, within, and between coefficients. The 12,985-child-year sample differs from the dated cognitive-only common sample. Smaller noncognitive and manual point estimates do not imply uniformly zero contributions.

Table B.8: Fixed O*NET five-domain contribution products

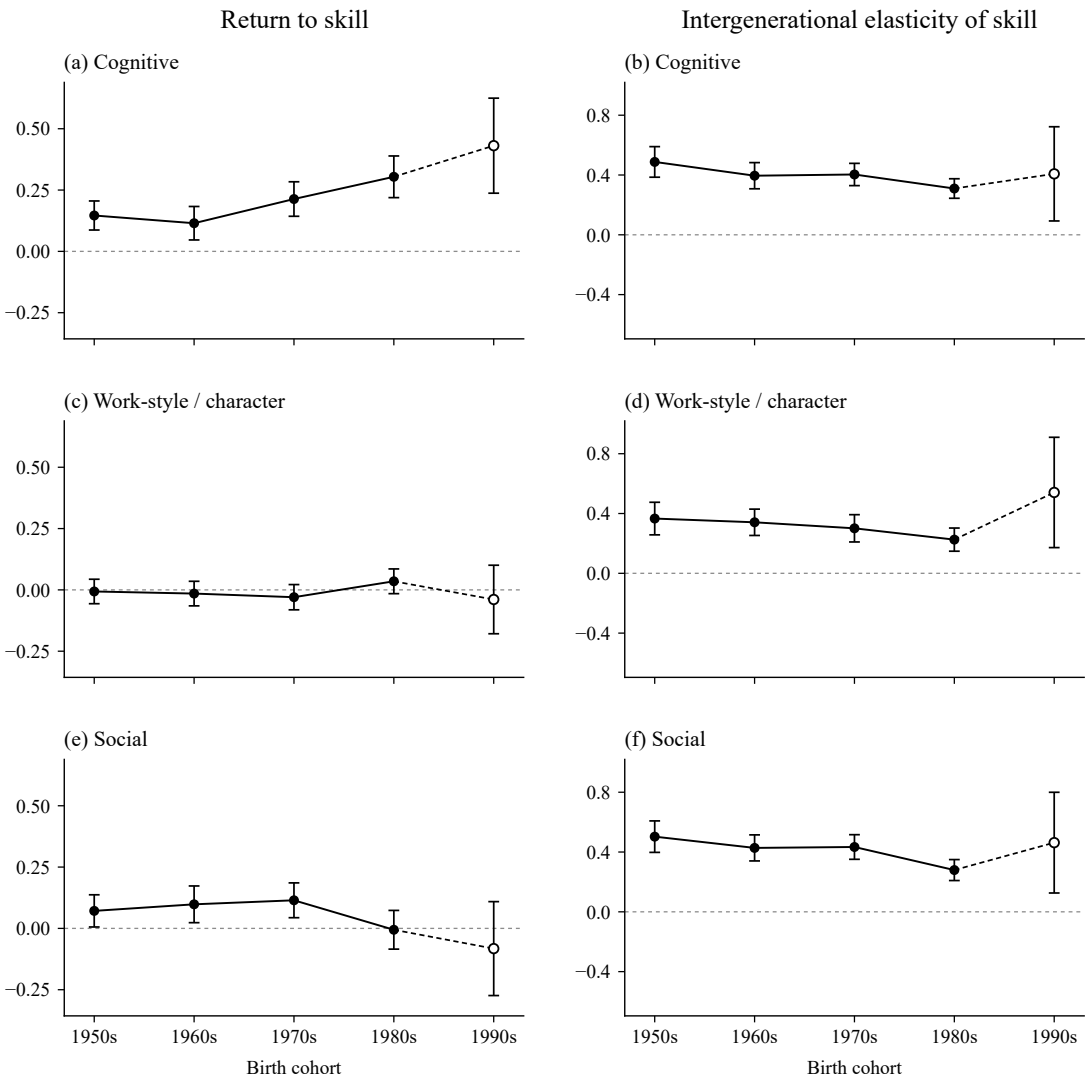
Cohort	Cognitive	Character	Social	Psychomotor	Physical
1950s	0.071 [0.040, 0.105]	-0.002 [-0.020, 0.016]	0.036 [0.004, 0.073]	-0.022 [-0.047, 0.005]	0.054 [0.023, 0.088]
1960s	0.045 [0.019, 0.078]	-0.005 [-0.022, 0.012]	0.042 [0.010, 0.078]	-0.002 [-0.024, 0.022]	0.020 [-0.001, 0.045]
1970s	0.086 [0.054, 0.119]	-0.009 [-0.025, 0.006]	0.050 [0.018, 0.084]	-0.051 [-0.077, -0.023]	0.077 [0.043, 0.113]
1980s	0.094 [0.060, 0.133]	0.008 [-0.003, 0.020]	-0.002 [-0.025, 0.021]	-0.026 [-0.048, -0.007]	0.045 [0.022, 0.073]
1990s†	0.176 [0.039, 0.383]	-0.021 [-0.114, 0.056]	-0.038 [-0.142, 0.052]	0.011 [-0.067, 0.096]	-0.000 [-0.050, 0.058]

Notes: Brackets report 95 percent whole-child bootstrap percentile intervals from 999 draws, conditional on constructed requirements. †: the 1990s age-33 target extrapolates from observed ages 30–32. Each entry is the corresponding joint income coefficient times its parental-income requirement gradient. All models retain the 13,738-child-year broad benchmark sample. Subdomain products need not sum to the product from a broad aggregate because the joint regressor space changes.

Table B.9: Full dated cognitive decomposition using 22 occupational families

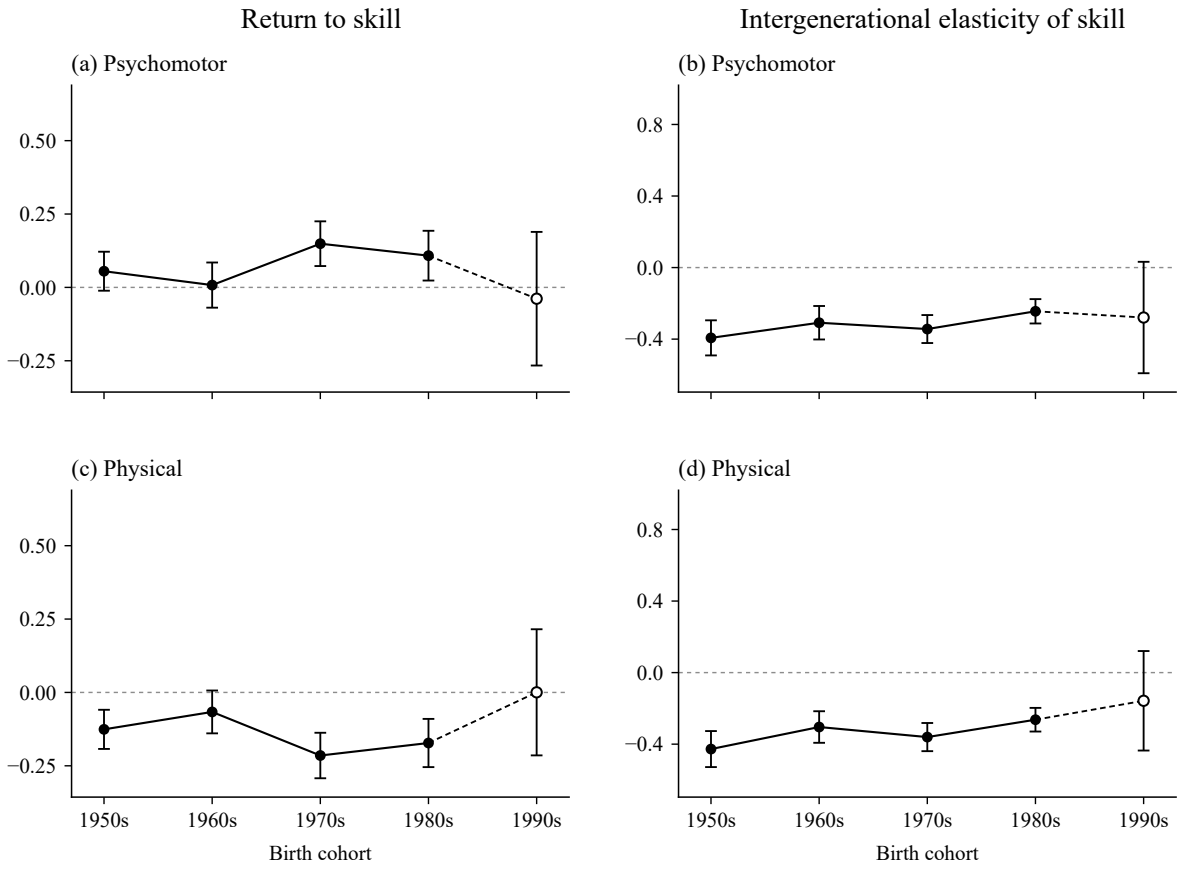
Cohort	IGE	Baseline	Reevaluation	Within	Between	Remainder
1950s	0.365 [0.286, 0.441]	0.108 [0.075, 0.145]	0.040 [0.012, 0.066]	-0.0038 [-0.0069, -0.0008]	-0.038 [-0.070, -0.006]	0.259 [0.184, 0.333]
1960s	0.367 [0.299, 0.444]	0.148 [0.107, 0.194]	-0.028 [-0.061, 0.001]	-0.0005 [-0.0015, 0.0005]	-0.058 [-0.083, -0.034]	0.306 [0.241, 0.377]
1970s	0.429 [0.354, 0.513]	0.116 [0.082, 0.158]	-0.012 [-0.042, 0.012]	0.0069 [0.0033, 0.0113]	-0.013 [-0.039, 0.015]	0.332 [0.260, 0.405]
1980s	0.323 [0.257, 0.383]	0.078 [0.046, 0.116]	0.038 [0.015, 0.056]	-0.0032 [-0.0097, 0.0023]	-0.028 [-0.067, 0.007]	0.239 [0.180, 0.298]
1990s†	0.613 [0.159, 1.029]	0.149 [0.054, 0.269]	0.118 [-0.002, 0.271]	-0.0059 [-0.0210, 0.0051]	-0.148 [-0.315, -0.001]	0.501 [0.068, 0.916]

Notes: Brackets report 95 percent whole-child bootstrap percentile intervals from 999 draws, conditional on constructed requirements. †: the 1990s age-33 target extrapolates from observed ages 30–32. Same 12,481 child-years and 4,024 children as detailed dated occupations. Raw descriptors are aggregated before frozen scoring, and returns are reestimated; only IGE must remain identical to the detailed model. Components add before rounding.



Notes: 95% child-clustered t intervals; all indices in this specification enter the income equation jointly.
Returns: log child family income per skill SD. Elasticities: skill SDs per log point of parental income.
Hollow markers: 1990s age-33 extrapolation (observed ages 30–32). Parent-income gradients evaluated at age 33.

Figure B.4: Cognitive, character, and social coefficients in the fixed five-domain model
Notes: Fixed broad O*NET profiles and the common 13,738-child-year five-domain model. All five indices enter the income equation jointly. Coefficient intervals use child-clustered t critical values; bootstrap intervals for products are in Table B.8. The youngest-cohort gradients extrapolate to age 33.



Notes: 95% child-clustered t intervals; all indices in this specification enter the income equation jointly.
Returns: log child family income per skill SD. Elasticities: skill SDs per log point of parental income.
Hollow markers: 1990s age-33 extrapolation (observed ages 30–32). Parent-income gradients evaluated at age 33.

Figure B.5: Psychomotor and physical coefficients in the fixed five-domain model

Notes: Fixed broad O*NET profiles and the common 13,738-child-year five-domain model. All five indices enter the income equation jointly. Coefficient intervals use child-clustered t critical values; bootstrap intervals for products are in Table B.8. The youngest-cohort gradients extrapolate to age 33.

Table B.10: Fixed-employment-weight profile sensitivity

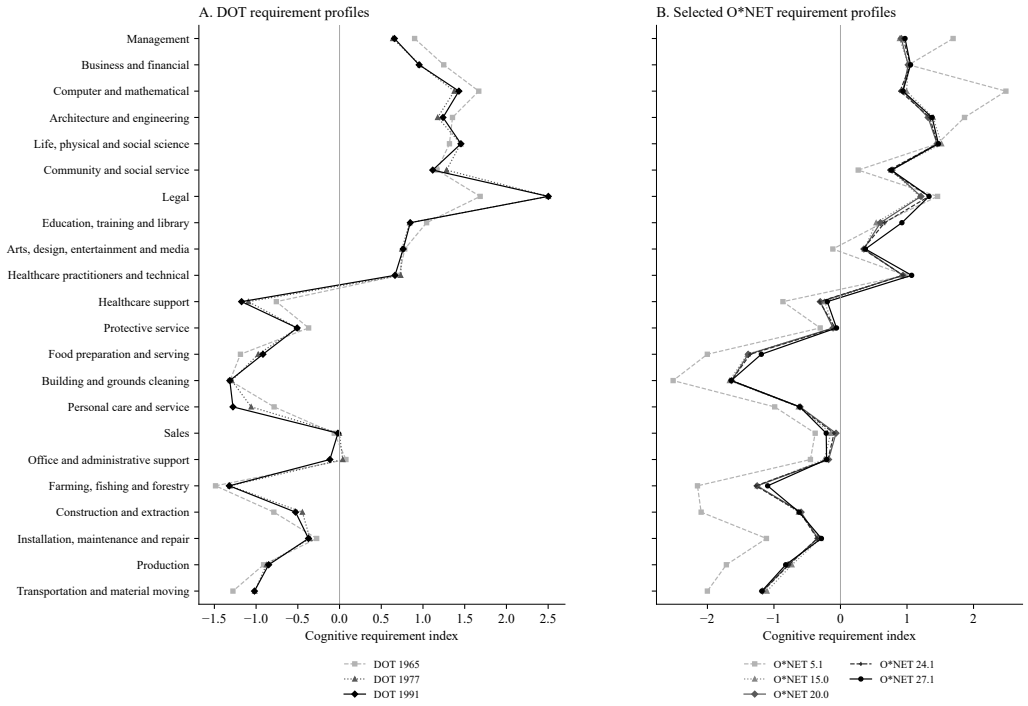
Period	Unit	N	Spearman	Δ mean	Rising (%)	Initial SD	Final SD
DOT 1965–1991	Detailed	267	0.930	-0.044	49.4	0.996	0.900
DOT 1965–1991	22 families	22	0.964	-0.037	44.5	0.872	0.773
O*NET 2003–2022	Detailed	392	0.886	0.417	72.5	1.521	0.972
O*NET 2003–2022	22 families	22	0.941	0.344	88.7	1.334	0.855
O*NET 2010–2022	Detailed	425	0.983	0.024	58.1	0.969	0.974
O*NET 2010–2022	22 families	22	0.988	0.051	72.2	0.833	0.855

Notes: Separate source-reference units. Fixed employment weights use original CPS reference totals for DOT and May 2018 OES totals for O*NET, held fixed across releases. Balanced-family constituent profiles restrict detailed occupation membership before raw-descriptor aggregation; they do not necessarily balance native specialties. These descriptive diagnostics were not substituted into income regressions.

Table B.11: Family profile sensitivity with balanced detailed constituents

Period	Unit	N	Spearman	Δ mean	Rising (%)	Initial SD	Final SD
DOT 1965–1991	22 families	22	0.964	-0.024	36.4	1.058	1.084
O*NET 2003–2022	22 families	22	0.944	0.341	77.3	1.432	0.950
O*NET 2010–2022	22 families	22	0.992	0.029	72.7	0.943	0.944

Notes: Separate source-reference units. Fixed employment weights use original CPS reference totals for DOT and May 2018 OES totals for O*NET, held fixed across releases. Balanced-family constituent profiles restrict detailed occupation membership before raw-descriptor aggregation; they do not necessarily balance native specialties. These descriptive diagnostics were not substituted into income regressions.



Notes: Earlier profile-line format with the same family ordering in both panels. DOT includes all three editions.
The O*NET panel selects five releases for readability; all 11 constructed releases appear in Figure 3 and Table 6.
Descriptors inherit the employment weights; fixed block PCA scoring is used without restandardizing family means. Source blocks have different rulers.

Figure B.6: Occupational family cognitive profiles across selected snapshots

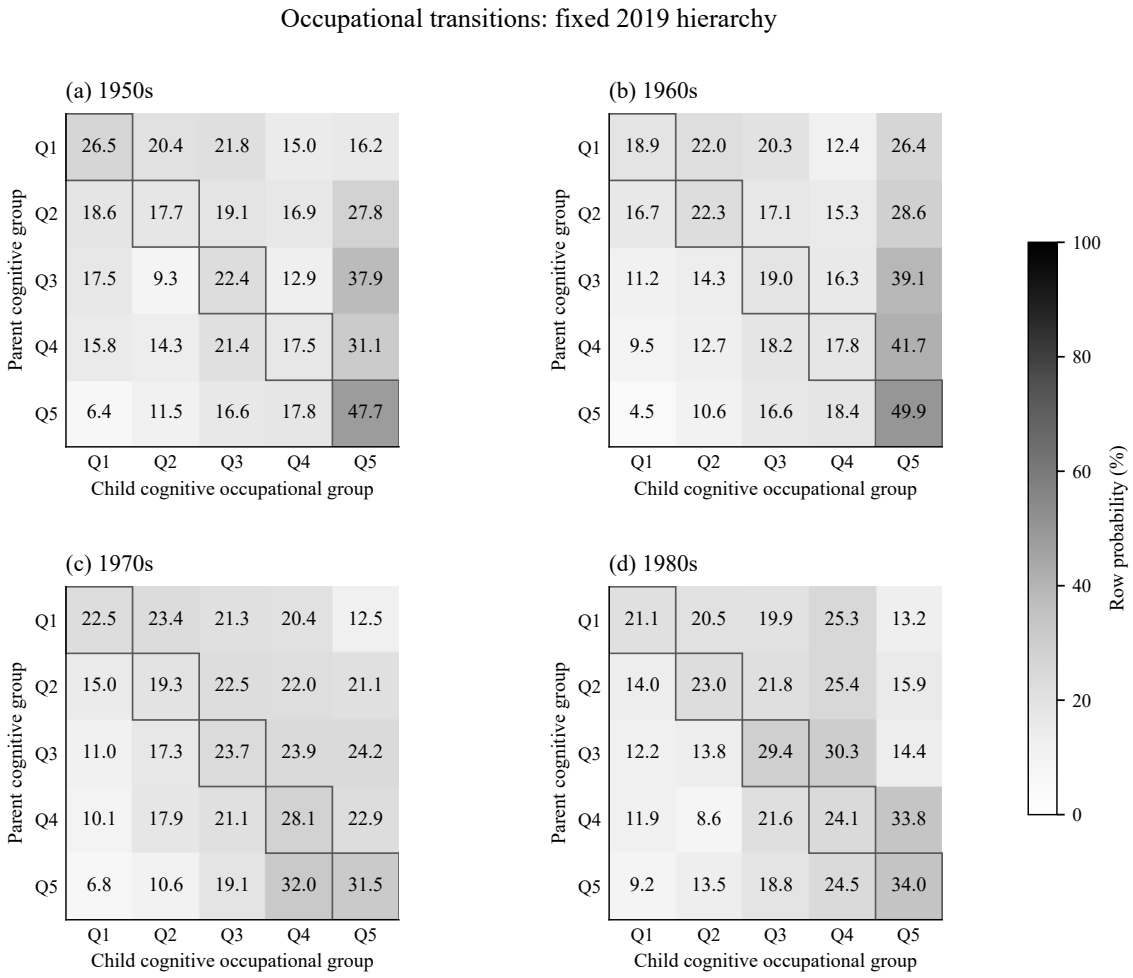
Notes: All 22 families use the same ordering. DOT shows all three editions; O*NET selects five releases for readability. The full snapshot heatmap remains in Figure A.3. Fixed source-block scoring permits changes within each block, but DOT and O*NET units are not cardinally calibrated.

C Occupational Transitions and Conditional Access

The transition sample is distinct from the dated cognitive sample. It uses the three-source seven-input cognitive profiles and requires both shared hierarchies for every contributing parent record. Fixed-profile occupation cutoffs are calculated on 448 reference occupations; child-period cutoffs use 267 DOT,423 O*NET 2010, or 448 O*NET 2019 reference occupations. A value on a cutoff enters the lower group; ties are never split. Continuous occupation percentiles use the number below plus half the number equal, divided by the reference count. Parental income quintiles use cohort-specific child-year inverse-CDF cutoffs, preserving income ties.

For origin standardization, the common pooled fixed-hierarchy reference shares for the 1950s–1980s are 16.90,15.84,14.47,17.64, and 35.16 percent across the five groups. These remain fixed in bootstrap draws. We resample whole children, recompute income quintile cutoffs, and use the same sampled children for the two hierarchy definitions. Skill profiles, occupation cutoffs, and parent-record weights remain fixed. For child observations assigned the latest historical source, including the entire provisional 1990s matrix group, the fixed and child-period hierarchies coincide by construction.

The conditional top-access statistic is $\Pr(g^o = 5 \mid g^p = a, z^p = z)$, where z^p is parental-income quintile. Its table is not row-stochastic. Of 250 origin-by-income-by-cohort cells across both hierarchies,102 have fewer than 30 distinct children and two are empty. The 30-child flag is descriptive rather than a power rule. Confidence intervals are withheld when fewer than 950 of 999 bootstrap draws have a valid denominator. Empty cells remain missing; observed zero probabilities are not given pseudocounts. No fitted margins-adjusted occupational association model has been estimated.



Cells show row probabilities (%), not counts. Q1 is lowest and Q5 highest; every row sums to 100.
Outlined diagonal: same cognitive group, not necessarily the same OCC2010 occupation.
Identical common-profile child-year sample across definitions; source cutoffs use equal-occupation references.

Figure C.1: Occupational transitions under the fixed cognitive hierarchy

Notes: Rows are parent occupational groups and columns are child groups under the same 2019 seven-input cognitive profile. The diagonal indicates the same cognitive group, not the same OCC 2010 occupation.

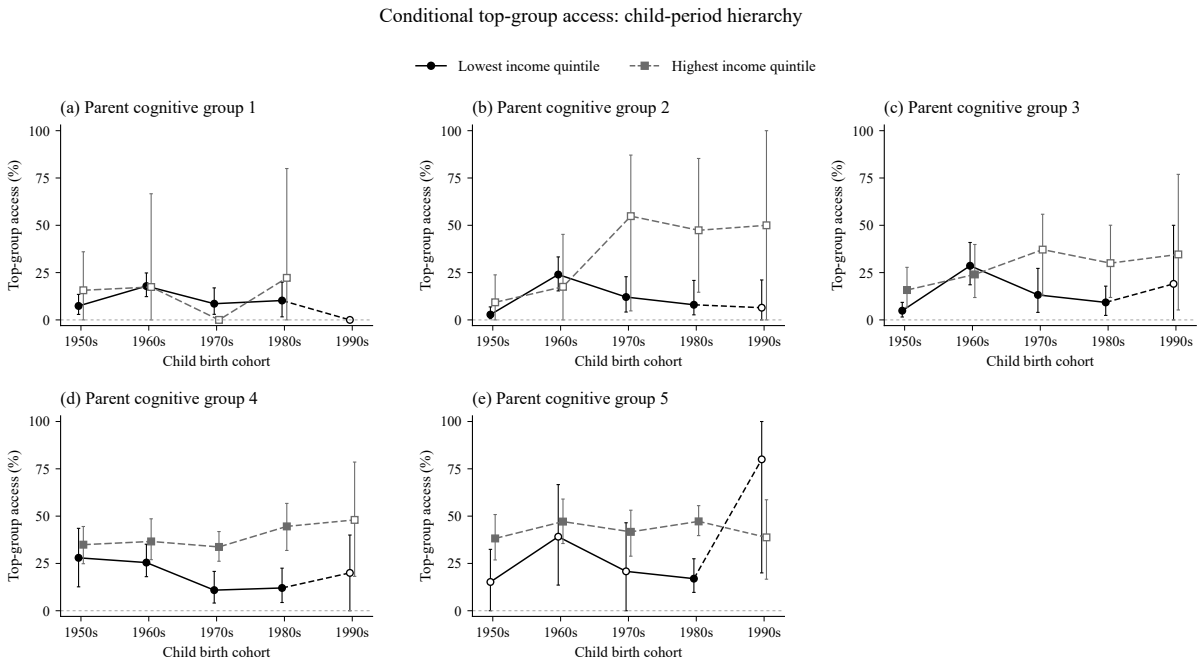


Figure C.2: Top cognitive-group access conditional on parental occupational origin

Notes: Child-period hierarchy. Cells condition jointly on parental occupational group and income quintile, so rows need not sum to one. Small conditioning samples are flagged in the original analysis; blank cells are not zero probabilities.

Table D.1: Three historical snapshots: joint three-domain labor-income decomposition

Cohort	IGE	Baseline	Reevaluation	Within	Between	Remainder
1950s	0.181 [0.111, 0.246]	0.181 [0.138, 0.232]	-0.059 [-0.110, -0.025]	0.0000 [0.0000, 0.0000]	-0.050 [-0.073, -0.026]	0.109 [0.048, 0.164]
1960s	0.217 [0.151, 0.284]	0.237 [0.191, 0.277]	-0.039 [-0.079, 0.001]	-0.0000 [-0.0006, 0.0006]	-0.108 [-0.138, -0.079]	0.127 [0.069, 0.187]
1970s	0.197 [0.134, 0.269]	0.165 [0.119, 0.214]	0.024 [-0.010, 0.054]	-0.0035 [-0.0081, 0.0001]	-0.086 [-0.126, -0.045]	0.098 [0.041, 0.160]
1980s	0.188 [0.133, 0.247]	0.185 [0.142, 0.236]	-0.023 [-0.058, 0.005]	-0.0011 [-0.0028, 0.0002]	-0.087 [-0.124, -0.054]	0.115 [0.069, 0.162]
1990s†	0.356 [0.057, 0.651]	0.192 [0.111, 0.327]	-0.028 [-0.130, 0.045]	0.0022 [-0.0027, 0.0101]	-0.095 [-0.240, 0.035]	0.285 [-0.049, 0.582]

Notes: Intervals in brackets are 95 percent whole-child bootstrap percentile intervals from 999 draws. All equations within a specification use identical observations and the specified controls. †: the 1990s group has observed ages 30–32; age-33 coefficients extrapolate. Components are separately projected on parental income with the same IGE controls; the remainder is obtained by subtraction. Rounding can affect displayed add-up.

D Labor-Income Results

Both generations use log real individual labor income. Child earnings assign household-head or spouse earnings according to the child's recorded relationship rather than automatically using head earnings. Parental labor income follows the saved father-first/mother-fallback labor-income rule. In the matched baseline,197 child-year observations select a different parent than the total-income design. We rebuild parent requirements from the labor-selected parent's corresponding earnings-window observations; year shares use all original earnings-window records.

The baseline retains 12,562 child-years and 4,065 children. The original total-income eligibility filters remain inherited from the starting analysis, so this is not an unrestricted earnings-observed population. The parent-age controls are also inherited from the original total-income window. In 749 matched rows, the actual mean age in the selected earnings window differs;197 rows change parent identity. These discrepancies are recorded but not corrected in the completed labor analysis. A future age-alignment sensitivity remains necessary before interpreting small total/labor differences as mechanisms.

The requirement system is the three-snapshot joint cognitive/noncognitive/manual model. The dated cognitive-only and 22-family extended analyses have not been repeated for labor income. Labor cognitive-only displays therefore show the cognitive components of a joint-skill model, not a newly estimated single-regressor model. Table D.1 retains the complete aggregate labor accounting. Through the 1980s, cognitive contributions are positive, but some noncognitive and manual net components are negative. For example, manual net contributions are approximately-0.028 in the 1970s and-0.009 in the 1980s, with intervals below zero. A cognitive-focused display should not erase those offsetting channels.

Table D.2: Total and labor income persistence on common child-years

Cohort	N	Children	Total IGE	Labor IGE	Paired difference
1950s	4661	1013	0.3634 [0.2883, 0.4403]	0.1722 [0.1052, 0.2406]	0.1912 [0.1157, 0.2685]
1960s	2730	875	0.3612 [0.2905, 0.4413]	0.2175 [0.1573, 0.2895]	0.1437 [0.0582, 0.2176]
1970s	2313	932	0.4542 [0.3823, 0.5364]	0.1973 [0.1215, 0.2677]	0.2568 [0.1828, 0.3398]
1980s	2479	1068	0.3184 [0.2622, 0.3840]	0.1875 [0.1375, 0.2444]	0.1309 [0.0716, 0.1921]
1990s†	278	229	0.5823 [0.1874, 0.9778]	0.3573 [0.0317, 0.6272]	0.2250 [-0.0830, 0.5103]

Notes: Parent income uses the corresponding total or labor definition in each model. These are common child-years of the three-snapshot joint-skill analyses, not the dated cognitive sample. Brackets are 95 percent child-cluster bootstrap intervals; differences use paired draws. Total income refers to family-unit money income, labor income to individual earnings. Their gap cannot be attributed solely to nonlabor income.

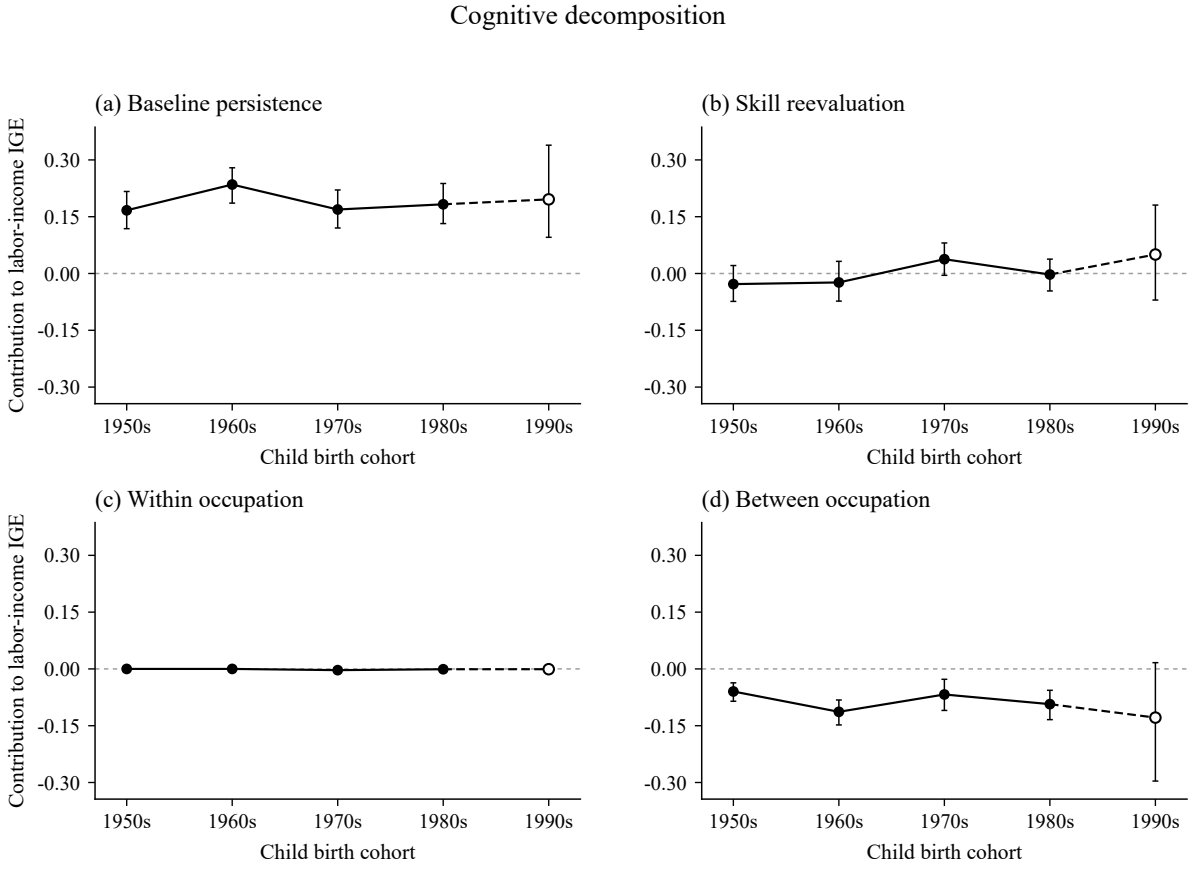


Figure D.1: Cognitive contributions to labor-income persistence

Notes: Both generations use individual labor income. Cognitive components are from the three-snapshot joint-skill model;12,562 child-years. This is not an extended dated cognitive-only regression.

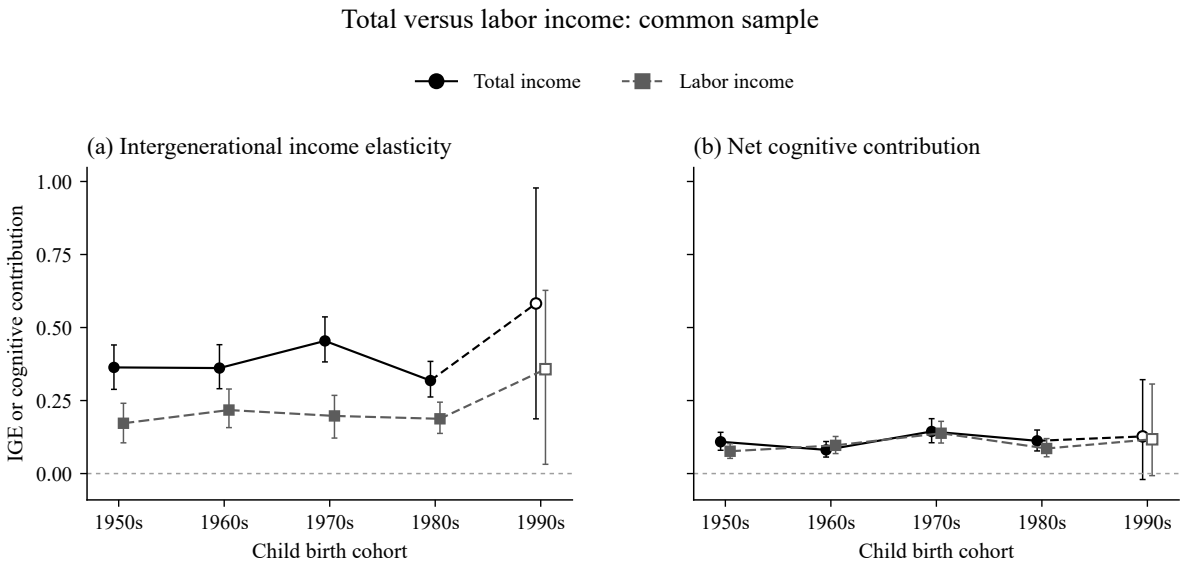


Figure D.2: Total-income and labor-income accounting on common child-years

Notes: Both generations use each outcome definition. Common child-year sample; parent selection follows the appropriate income-specific rule. Different recipient units and income content preclude interpreting the gap as a pure nonlabor-income effect.

E Skills Re-evaluation in Labor Market

Now we show the proof of Proposition 1.

Proof.

$$\beta^{IGE} = \frac{cov(y^o, y^p)}{var(y^p)} = \frac{cov(\tau_c^o c^o + \tau_n^o n^o + \varepsilon^o, \tau_c^p c^p + \tau_n^p n^p + \varepsilon^p)}{var(y^p)}$$

$$= \frac{\tau_c^o \tau_c^p cov(c^o, c^p) + \tau_n^o \tau_n^p cov(n^o, n^p) + \tau_c^o \tau_n^p cov(c^o, n^p) + \tau_n^o \tau_c^p cov(n^o, c^p)}{var(y^p)}$$

$$= \frac{\tau_c^o \tau_c^p \rho_c + \tau_n^o \tau_n^p \rho_n + \tau_c^o \tau_n^p \rho_c \gamma + \tau_n^o \tau_c^p \rho_n \gamma}{var(y^p)}$$

$$= \frac{(\tau_c^p + a)\tau_c^p \rho_c + (\tau_n^p + b)\tau_n^p \rho_n + (\tau_c^p + a)\tau_n^p \rho_c \gamma + (\tau_n^p + b)\tau_c^p \rho_n \gamma}{var(y^p)}$$

$$= \frac{(\tau_c^p)^2 \rho_c + (\tau_n^p)^2 \rho_n + \gamma \tau_c^p \tau_n^p (\rho_c + \rho_n) + a\rho_c(\tau_c^p + \gamma \tau_n^p) + b\rho_n(\tau_n^p + \gamma \tau_c^p)}{var(y^p)}$$

If no skills re-evaluation is present, then $\tilde{\beta}^{IGE} = \beta^{IGE}$ holds trivially.

E.1 Comparative Statics

If technological change shifts the returns to skills, then

$$\tilde{\beta}^{IGE} - \beta^{IGE} = \frac{a\rho_c(\tau_c^p + \gamma \tau_n^p) + b\rho_n(\tau_n^p + \gamma \tau_c^p)}{var(y^p)}$$

If technological change decrease the returns to one type of skill (e.g. $a < 0$), intergenerational mobility will increase (intergenerational persistence will decrease), so long as

$$a < -b \frac{\rho_n}{\rho_c} \frac{\tau_n^p + \gamma \tau_c^p}{\tau_c^p + \gamma \tau_n^p}$$

where it depends on the magnitude of the original returns to skills and the intergenerational persistence of skill supply.

Therefore, the effect of technological change on intergenerational mobility is ambiguous. □