

The Direct and Indirect Effects of Genetics and Education

Senan Hogan-Hennessy*

Economics Department, Cornell University[†]

First draft: 12 March 2026

This version: 18 July 2026

Abstract

Genes associated with educational attainment causally improve labour market income, but the economic mechanism behind this relationship is not clear. Using quasi-random variation in genetic inheritance across siblings in the UK Biobank, I estimate the causal effect of the Education PolyGenic Index (Ed PGI) on education and income. A one standard deviation increase in the Ed PGI raises completed education by 0.5 years and later-life income by around 5 percent (replicating the main estimates in [Carvalho 2025](#)). I then decompose this total genetic income effect into an indirect channel operating through education years and a residual direct effect, using a causal mediation framework. Unlike structural model-based decompositions, this approach is design-based; one remaining source of uncertainty (the causal return to an extra year of education) is handled transparently through a sensitivity analysis. At correlational education returns of around 6 percent, 65 to 75 percent of the total genetic income effect operates through the years-of-education channel. Quasi-experimental estimates from the economics literature for Britain imply higher returns to education, suggesting that the mediated share is larger and that the majority of the Ed PGI's income effect operates through years of completed education.

Keywords: Education, Mediation, Polygenic indices, Nature/nurture.

JEL Codes: D31, D91, I24, J24, Z00.

*This work has been supported by research grants from the Centre for the Study of Inequality and the Industrial Labour Relations School, Cornell University, conducted under UK Biobank Project #335854. Thanks to Harry Patrinos for providing the collected literature review estimates (from [Patrinos & Psacharopoulos 2026](#)). For helpful comments I thank Neil Cholli, Ryan Dycus, Benny Goldman, Takuma Habu, Jake Meyer, Douglas Miller, Lauri Kytömaa, Zhuan Pei, Evan Riehl, and David Slichter. Any comments or suggestions may be sent to me at seh325@cornell.edu.

[†]Address: Uris Hall #447, Economics Department, Cornell University NY 14853 USA.

The difference of natural talents in different men, is, in reality, much less than we are aware of [...]. The difference between the most dissimilar characters, between a philosopher and a common street porter, seems to arise not so much from nature, as from habit, custom, and education.

— [Smith \(1776\)](#), *The Wealth of Nations*, Volume 1 p. 16.

1 Introduction

Genes associated with educational attainment causally improve labour market outcomes, but the mechanism behind this relationship remains poorly understood. Individuals with more education-associated genes have higher incomes and more years of education, yet it is unclear how much of their income advantage reflects the additional education their genes induce and how much operates through other pathways. Understanding the relative importance of these channels provides insight into the process of human capital formation. If most of the income effect operates through years of completed education, then education-linked genetic variation influences labour market outcomes primarily by altering schooling decisions and educational investments, rather than through pathways outside completed years of formal education.

The Education PolyGenic Index (Ed PGI) is a weighted sum of thousands of genetic variants that together explain around 12 to 15 percent of the variation in educational attainment ([Okbay et al. 2022](#)); individuals with a higher Ed PGI value have more years of completed education, and higher occupation-measured incomes — and this relationship is causal ([Carvalho 2025](#)). This paper asks how much of the Ed PGI’s causal effect on income operates through completed education years, and how much remains outside that channel.

This paper makes two contributions. First, it provides an explicit quantitative decomposition of the Ed PGI’s causal effect on labour-market income, separating the share operating through completed education years from a residual direct channel. Existing research establishes that the Ed PGI affects both educational attainment and labour-market outcomes, but does not quantify the relative importance of these two pathways. Second, the paper develops a way to conduct this decomposition when education years are not randomly assigned. The analysis combines the causal effect of the Ed PGI on education with estimates of the return to an additional year of schooling, using the correlational return as an initial benchmark and quasi-experimental estimates from the British literature to guide a sensitivity analysis. This approach makes the central source of uncertainty explicit and shows how conclusions about the two channels change with the assumed return to education, without requiring a structural model of schooling or labour-market choices.

I use data regarding 24,743 individuals from the UK Biobank to estimate the causal effect of the Ed PGI on education and labour market income, following the design of [Carvalho \(2025\)](#). The Ed PGI varies quasi-randomly across individuals: when a parent carries two different copies of a genetic variant, which copy the child inherits is determined by a random biological draw. I control for imputed parental Ed PGI values, constructed from the genetic data of siblings also enrolled in the UK Biobank ([Young et al. 2022](#)), to isolate this random inheritance component. A one standard deviation increase in the Ed PGI increases years of education completed by 0.5 years and raises later life income by around 5 percent, confirming the main findings in [Carvalho \(2025\)](#).

The first part of the analysis establishes the two causal effects required for the decomposition: the effect of the Ed PGI on education and its total effect on labour-market income. The research design and these total-effect estimates build on [Carvalho \(2025\)](#), but they provide the necessary empirical foundation for the paper’s causal mediation analysis. After establishing these inputs, the second part quantifies how much of the income effect operates through completed education years.

I apply a causal mediation framework to quantify how much of the Ed PGI’s causal effect on income operates through completed education years. The decomposition combines the Ed PGI’s average causal effect on education years with the return to an additional year of education. The latter is not identified by the genetic design, so I first use the correlational return as a benchmark, and then evaluate the decomposition over the distribution of quasi-experimental returns estimated in the British economics literature ([Patrinos & Psacharopoulos 2026](#)). This approach makes the remaining source of uncertainty transparent without requiring a structural model of education or labour-market choices.

My main finding is that the Ed PGI raises labour market income primarily through education years. At correlational returns to education of around 6 percent, roughly 65 to 75 percent of the total genetic income effect is mediated through the education years channel. The sensitivity analysis then evaluates this decomposition over the distribution of returns to education estimated in the British economics literature. Across this range, the mediated share rises further as the return to education increases, implying that the majority of the Ed PGI’s labour market effect operates through educational attainment rather than residual non-years-of-education pathways.

Modern genetic data provide a new way to study a central question in the economics of human capital formation: how initial endowments interact with environments and investments to shape education and later-life outcomes ([Cunha & Heckman 2007](#), [Beauchamp et al. 2011](#),

[Biroli et al. 2025](#)). This paper contributes to a growing literature applying these data to socioeconomic outcomes. [Papageorge & Thom \(2020\)](#), [Bryson et al. \(2025\)](#) establish that the Ed PGI predicts labour-market outcomes, although [Schork et al. \(2022\)](#) shows that these associations may partly reflect parental genetic channels rather than causal effects of an individual’s own genes. Using the sibling-imputation procedure developed by [Young et al. \(2022\)](#), [Carvalho \(2025\)](#) addresses this concern by controlling for imputed parental PGIs and estimating causal effects; [Muslimova et al. \(2025\)](#), [Sanz-de Galdeano & Terskaya \(2025\)](#) exploit related within-family variation to study complementary questions. A parallel literature uses structural models of education and labour-market choices to examine how the Ed PGI shapes socioeconomic outcomes ([Barth et al. 2022](#), [Houmark et al. 2024](#), [Bolt et al. 2025](#)). Together, these studies establish that the Ed PGI predicts and causally affects socioeconomic outcomes, and structural models provide interpretations of the choices underlying these relationships. They do not, however, provide a design-based quantitative decomposition of the income effect into the component operating through completed education years and the component remaining outside that channel.

Causal mediation decomposes an average treatment effect into an observed mechanism and a remaining direct effect, where credible identification requires quasi-random variation in the mediating channel ([Imai et al. 2010](#)) — see [Huber \(2020\)](#) for an overview of the causal mediation literature.¹ When the mediating channel is not randomly assigned (as education years plainly are not), conventional mediation estimates are not causal and suffer from selection bias ([Hogan-Hennessy 2025](#)). I instead combine the identified causal effect of the Ed PGI on education with estimates of the return to an additional year of schooling. The correlational return provides an initial benchmark, while a sensitivity analysis shows how the decomposition changes with different values for causal returns to education. I discipline this analysis using the distribution of quasi-experimental estimates from the economics literature for Britain. The resulting design-based decomposition quantifies the contribution of education years without requiring a structural model of schooling or labour-market choices.

This paper proceeds as follows. [Section 2](#) introduces the UK Biobank data and the construction of the Ed PGI. [Section 3](#) and [Section 4](#) present the framework and empirical results, respectively, of the two inputs to the causal mediation decomposition: the Ed PGI’s effect on education and its total effect on labour-market income. [Section 5](#) develops the framework for separating the education-years and residual channels, and [Section 6](#) presents the empirical results, and a brief discussion. [Section 7](#) concludes.

¹Analyses of mechanisms in applied economics typically rely on suggestive evidence instead, which does not identify or quantify the role of a mechanism. See [Blackwell et al. \(2024\)](#), [Green et al. \(2010\)](#).

2 Data and Genetic Context

This paper draws on data from the UK Biobank, a large-scale cohort study of the British population with both rich socioeconomic measures and detailed genetic data. I construct the Education PolyGenic Index (Ed PGI) summarising each participant’s genetic predisposition toward educational attainment, and impute parental PGI values using genetic data from siblings. Together, these data enable a within-family research design that separates the role of inherited genetic factors from the shared family environment in shaping socioeconomic outcomes.

2.1 UK Biobank (UKB)

The UK Biobank (UKB) is a large-scale prospective cohort study comprising approximately half a million members of the British public, recruited between 2006 and 2010. Letters were sent to a random sample of addresses in the National Health Service registry, and the final UKB data represent voluntary respondents. Participants, aged 40-69 years at recruitment, underwent comprehensive baseline assessments, were asked many questions for socioeconomic information either by interview, automated survey answers, or follow-up sessions. Key socioeconomic variables include educational attainment (measured in education qualifications), household income brackets, employment status, and industry of employment. See [Sudlow et al. \(2015\)](#) for a full description of the UKB organisation and data.

The UKB collected detailed information on educational qualifications, including completion of primary school, GCSEs (typically completed at age 16), A-Levels (completed at age 18), and completion of an higher education degree. I convert these qualifications into education years based on each participant’s highest reported qualification, following the International Standard Classification of Education.² While the UKB collected household income data, this information was recorded only in categorical bins; I assign midpoint imputed values and update to 2024 GBP for a coarse measure of household income. To accompany these measures with a finer measure of income, I also calculate occupation income. I construct an individual-level income measure using participants’ 4-digit Standard Occupational Classification (SOC) codes. Specifically, I impute occupation-based wages by estimating hourly earnings for each SOC category using data from the British Annual Survey of Hours and

²British vocational degrees are inconsistently coded by the International Standard Classification of Education. I follow [Okbay et al. \(2022\)](#), and assign education years as the age they reported leaving full-time education minus five, among those with a vocational degree as their highest qualification.


Table 1: Descriptive Statistics, UKB.

	Analysis sample		Entire sample	
	Mean	SD	Mean	SD
<i>Demographics:</i>				
Male	0.44	0.50	0.46	0.50
Age	55.04	7.05	56.55	8.09
Lives in city	0.86	0.35	0.85	0.35
Ethnicity = European	1.00	0.00	0.82	0.38
Any siblings also in UKB?	1.00	0.00	0.08	0.28
Count of siblings in UKB	1.09	0.32	0.09	0.32
<i>Genetic Measures:</i>				
Ed PGI	0.04	1.00	0.00	1.00
Ed PGI, imputed parental mean	-0.02	0.75		
Other PGI: ADHD	-0.02	1.00	-0.00	1.00
Other PGI: Asthma	-0.01	1.00	0.00	1.00
Other PGI: Bipolar	-0.01	1.00	-0.00	1.00
Other PGI: BMI	-0.02	0.99	0.00	1.00
Other PGI: Diabetes (type 2)	-0.02	1.00	0.00	1.00
Other PGI: Height	0.00	1.01	-0.00	1.00
Other PGI: Schizophrenia	-0.04	1.00	0.00	1.00
<i>Education:</i>				
Education years	13.84	3.36	13.72	3.48
Age left education	16.67	2.47	16.51	3.04
Qualification, University degree	0.33	0.47	0.33	0.47
Qualification, A-Levels	0.46	0.50	0.45	0.50
Qualification, GCSEs	0.76	0.43	0.71	0.45
Qualification, Professional degree	0.32	0.47	0.29	0.45
Qualification, Vocational degree	0.23	0.42	0.19	0.39
Qualification, No official qualifications	0.13	0.34	0.17	0.38
<i>Labour Market Outcomes:</i>				
Occupation hourly wage, £	21.36	9.15	22.02	9.49
Occupation annual income, thousands £	30.27	19.02	31.95	20.04
Average hours worked, per week	34.73	12.58	35.31	12.70
Household income, < £18k	0.13	0.33	0.20	0.40
Household income, £18 – 31k	0.23	0.42	0.22	0.41
Household income, £31 – 52k	0.28	0.45	0.22	0.42
Household income, £52 – 100k	0.22	0.41	0.17	0.38
Household income, £100k <	0.05	0.22	0.05	0.21
Household income, thousands £ (midpoint imputed)	67.13	40.07	60.95	41.46
Observations	24,743		496,004	

Note: This table shows the Mean and Standard Deviation (SD) of UKB variables, comparing the analysis subsample and the entire UKB sample. The analysis sample refers to individuals with at least one family member also included in the UKB (sibling or parent), have non-missing education and labour market occupation codes, and are of European ancestry. All genetic variables are coded as missing among those who report ancestry that is not European.

Earnings, conditional on age, gender, and year.³ This approach provides a more granular and continuous measure of economic status than the categorical household income variable. I create the annual occupation income by multiplying this by the number of hours they reported working, times 52 weeks in the year. This gives 496,004 individuals in the UKB, with non-missing data in relevant variables.

Figure 1: Distribution of UKB Education Measures.



Note: The left plot shows the observation count in the UKB analysis sample with each number of reported education years, coded from education qualifications. The right plot shows the mean occupation imputed annual income for each level of education years, with a line of best fit calculated from a linear regression (with annual income measured in a log specification); the slope refers to the estimated coefficient from this linear regression.

For the purposes of this study, the primary analysis sample is restricted to those with at least one family member, sibling or parent,⁴ also enrolled in the UKB, and non-missing data in occupation code plus education qualifications. This gives 24,743 individuals in the UKB. The presence of multiple family members within the dataset provides a unique opportunity to control for shared genetic factors that might otherwise bias estimates in standard cross-sectional analyses.

³This approach uses replication data files provided by [Kweon et al. \(2025\)](#), also used by [Carvalho \(2025\)](#); see [Appendix A.1.2](#).

⁴The UKB does not personally identify participants, nor their family members, so I follow the genetic literature by identifying family clusters by measures of genetic similarity. [Appendix A.1.1](#) further describes the family definitions in the UKB.

The analysis sample is also restricted to individuals who report to have primarily European ancestry. The genetic measures used in this paper have generally been constructed from primarily white, European descended populations, and have little predictive power out of sample, so are only relevant in this subsample. It is an unfortunate fact that systematic genetic measures, such as those considered in this paper, have only been constructed from Biobank samples where the vast majority are white and/or European descended; see [Martin et al. \(2017\)](#) for a further discussion. A limitation of this study is that its results only apply to this subset of the UKB population.

The only notable differences between the analysis subsample and the entire UKB participant sample is that the analysis sample has more siblings, and is entirely European descended; no other variables are significantly different for the analysis subsample. [Table 1](#) summarises the UKB data, for the analysis sample and among the entire UKB sample. [Figure 1](#) summarises the correlation between education years and occupation coded income.

2.2 Education PolyGenic Index (Ed PGI)

UKB participants also provided saliva samples that were subsequently genotyped, producing standardised genetic data files for statistical analysis of the relationship between genetics and outcomes.

While genetic data provide rich information about individual variation, the relationship between specific genes and social outcomes like education is extraordinarily complex. The biological pathways from DNA to observable traits remain largely opaque — the exact relationship between a single gene, its resulting protein (or its moderating role), and the resulting outcomes are only known for a small number of genes in the human genome. For example, a single gene (and its recessive variant) is known to lead to a protein which does not break down the by-products of alcohol digestion; individuals with this genetic variant generally do not drink alcohol at all because of the biological limit on digestion ([Millwood et al. 2023](#)). Complex social behaviours, on the other hand, have no single gene with a known biological pathway, so cannot be summarised with one single gene, nor single biological pathway. Instead, complex social behaviours are a result on thousands of genes, and their interactions with the environment — they are polygenic.

Rather than attempting to trace all of these intricate biological mechanisms, researchers have developed PolyGenic Indices (PGIs) as dimensional reduction tools that summarise the statistical relationship between thousands of genetic variants and outcomes of interest. This

statistical aggregation is particularly valuable because individual genetic variants typically have tiny associations with social outcomes like educational attainment — often explaining less than 0.01% of the variation. However, by combining information across thousands of genetic variants, each contributing a small predictive association, PGIs can capture meaningful predictive power of genes for some social outcomes. These indices are fundamentally correlational measures: they capture statistical associations between genetic variants and outcomes observed in large samples, not well understood biological pathways.

The Education PGI (Ed PGI) is the PolyGenic Index built to predict years of completed education, calculated from weights derived from a Genome-Wide Association Study (GWAS), which analysed the correlation between genetic variants and education years. In a GWAS, researchers test the correlation between each Single Nucleotide Polymorphism (SNP) — a one-letter change in the DNA sequence at a specific location — and the outcome of interest (Uffelmann et al. 2021). I use the weights calculated in the GWAS conducted by Okbay et al. (2022), which used a sample of over 3 million individuals of European ancestry, primarily from 23&Me data.


The Ed PGI is defined as the weighted sum of whether individual i has the relevant SNP at each DNA location j ,

$$Z_i = \sum_{j=1}^J w_j x_{i,j}$$

where w_j are the Ed PGI weights (provided in the supplementary data to Okbay et al. 2022), and $x_{i,j}$ is an indicator of whether individual i has the relevant SNP at DNA location j .⁵ The Okbay et al. (2022) GWAS found that around 4,000 individual SNPs were significant in predicting education years, after passing a battery of biostatistical tests and finding no evidence for significant interactions. The resulting Ed PGI is standardised to have mean zero and standard deviation one within the entire sample, making it interpretable as the number of standard deviations from the population mean. The Ed PGI captures approximately 12–15% of the variation in educational attainment, making it one of the most predictive PGIs for any complex behavioural trait to date.

If education had no relationship with SNP variants, then we would expect the correlation between the Ed PGI and education years was indistinguishable from zero; it is, however, highly correlated with education years in the UKB analysis sample. A one standard deviation

⁵Every person has two copies of each chromosome, so that $x_{i,j}$ takes values 0, 1, 2 corresponding to how many copies of the SNP individual i has inherited. The genetics literature has generally accepted that complex social behaviours are explained by additive effects (see Hill et al. 2008), and Okbay et al. (2022) found little evidence for dominance in the Ed PGI SNPs, so that the Ed PGI makes sense as a weighted sum.

Figure 2: Distribution of Education PolyGenic Index (Ed PGI).

Note: The left plot shows a kernel density estimate for the distribution of the Ed PGI among the UKB analysis sample. Ed PGI is normalised to have mean 0 and standard deviation 1 among the entire UKB sample; it is roughly normally distributed regardless of this normalisation. The right plot shows a bin scatter, visualising a (semi-parametric) correlation between the Ed PGI and education years. The slope refers to the estimated coefficient from a linear regression between the Ed PGI and education years.

increase in the Ed PGI is associated with 0.93 extra education years (standard error 0.02), with an R^2 value of 0.10 for the raw regression of education years on Ed PGI. It is roughly normally distributed, with standard deviation normalised to 1. This distribution is not surprising, as the Ed PGI is the weighted sum of a few thousand binary variables, so we should expect convergence towards a normal distribution among the population. The main take-away is that the correlation between a summary measure of genetic values is strong, and has meaningful predictive power for education years.

The Ed PGI is not the only PGI that can be constructed. Following the same methodology, I construct additional PGIs for six health-related traits that may confound the relationship between genetics and socioeconomic outcomes: asthma, bipolar disorder, Body Mass Index (BMI), height, schizophrenia, and type 2 diabetes. Each PGI is calculated using the same approach as the Ed PGI — taking weights from a published GWAS, conducted on independent populations, and computing weighted sums of the relevant genetic variants in the UKB data. These other PGIs serve as control variables in my analysis for two reasons. First, genetic


variants often exhibit pleiotropy, meaning a single variant can influence multiple traits simultaneously. For instance, genetic variants associated with education might also affect other factors (e.g., mental health conditions), which could independently influence labour market outcomes. Second, systematic differences in gene frequencies across ancestry groups can create spurious correlations between any PGI and outcomes if health conditions vary across these groups for unrelated reasons.

2.3 Imputing Parental PGIs from Siblings

It would be ideal if the UKB contained genetic information for the parents of every UKB participant. While there are 6,239 participants with a parent also included in the UKB, there are only 72 with both parents. To overcome this missing data problem, I also use siblings' genetic data to impute the mean parental PGIs from the sample of 24,743 participants with at least one family member (parent or sibling) also in the UKB — the analysis sample described above. The majority of this analysis sample have genetic data on one sibling member, so I refer to this as sibling imputation.

I use the [Young et al. \(2022\)](#) procedure to impute mean parental PGI values using genetic data on participants' siblings. The key insight is that since siblings both inherit their DNA from the same two parents, observing two siblings' genotypes allows us to partially reconstruct the parental genotypes they were drawn from. At each SNP, the imputation proceeds in two steps. First, the procedure classifies each sibling pair as being in one of three identity-by-descent (IBD) states using a hidden Markov model that exploits information across all SNPs: IBD0, in which the siblings share no alleles inherited from the same parent; IBD1, in which they share one such allele; and IBD2, in which they share both. This IBD classification has been shown to be correct 99.65% of the time in validation samples ([Young et al. 2022](#)). Second, given the inferred IBD state, the procedure imputes the sum of the two parental genotypes at that SNP — which is the quantity needed to construct a parental PGI — by summing the directly observed parental alleles and replacing any unobserved alleles with the population allele frequency $f \in [0, 1]$.

[Figure 3](#) illustrates this logic for the case where both the child and their sibling have observed SNP value of 1 (one copy of the reference allele). In IBD0, the siblings share no alleles, so all four parental alleles are directly observed across the two children, and the parental sum is known exactly as 2. In IBD1, one allele is shared between the siblings, leaving one parental allele unobserved; depending on whether the shared allele is the reference allele or not, the imputed parental sum is either $1 + f$ or $2 + f$, depending on value of the inherited

Figure 3: Imputing One SNP for Parents from Sibling Data.

Note: This figure illustrates the logic of imputing the sum of parental genotypes from sibling data at a single SNP, following [Young et al. \(2022\)](#). At each SNP, individuals have a genotype of 0, 1, or 2, representing the number of copies of a reference allele inherited from their parents; this figure shows the case where both the child and their sibling have genotype 1, across the three possible identity-by-descent (IBD) states. In IBD0, the siblings share no alleles by descent, so all four parental alleles are directly observed and the parental sum is known exactly as 2. In IBD1, the siblings share one allele by descent, leaving one parental allele unobserved; the imputed parental sum is $1 + f$ or $2 + f$, depending on whether the shared allele is the reference allele, where f denotes the population frequency of the reference allele. In IBD2, the siblings share both alleles by descent, so only two parental alleles are observed and the imputed parental sum is $1 + 2f$. In all cases, unobserved parental alleles are replaced by the population allele frequency f .

shared allele. In IBD2, both siblings inherited the same two alleles, so only two of the four parental alleles are observed (summing to 1), and the two unobserved alleles are each imputed as f , giving sum of parents $1 + 2f$. In all cases, the imputed value is the conditional expectation of the parental sum given the observed sibling genotypes and IBD state, where unobserved alleles are replaced by the population frequency. Averaging across the three IBD

states — which occur with probabilities $1/4$, $1/2$, and $1/4$ respectively — siblings jointly reveal on average three of the four parental alleles, meaning approximately 75% of parental genetic information is directly observed rather than imputed. See [Young et al. \(2022\)](#) for further reasoning on the imputation performance relative to observed parents in the UKB.

2.4 Ethical Considerations

The genetic measures in this study should not be misinterpreted as evidence of genetic determinism. Having a higher Ed PGI does not destine someone to higher education or income; it represents a statistical tendency, not fate. These tendencies are fundamentally shaped by social policies, institutional structures, and cultural contexts that vary across time and place. Just as eyesight correcting glasses can eliminate the practical consequences of genetic variation in eyesight, education and economic policies can modify or eliminate the consequences of genetic variation in traits relevant to socioeconomic outcomes.

Furthermore, this research provides no support for ideas about “natural” social hierarchies or genetic justifications for inequality. In addition, genetic associations identified within one ancestry group cannot be meaningfully compared across groups due to different genetic architectures, environmental contexts, and historical experiences. Any claim using the measures considered in this paper for inference across racial groups would be both invalid and ethically deficient. The appropriate interpretation of these results is that they highlight how both social and genetic factors influence socioeconomic outcomes, which should inspire greater solidarity and support for policies that promote equality of opportunity rather than acceptance of inequality as “natural” or inevitable.

3 Random Variation in Genetic Inheritance

Genetic variants are not randomly assigned, which means that simple correlations between the Ed PGI and observed outcomes cannot straightforwardly be interpreted as causal. Genetic variation arises from inheritance: at conception, individuals inherit 23 chromosomes from each parent, obtaining half of their genetic variants from their mother and half from their father. Those who inherit a higher Ed PGI therefore typically also inherit familial environments shaped by parents who themselves possess a higher Ed PGI. Consequently, correlations between an individual’s Ed PGI and outcomes may reflect effects operating through the family environment. However, inheritance also contains a random component: conditional on

parental genetics, Mendelian inheritance generates random variation in the genetic variants a child receives. I use this variation to estimate the causal effects of the Ed PGI.

This research design establishes the two causal effects used in the later mediation analysis: the Ed PGI’s effect on education and its total effect on labour-market income. The design and resulting estimates build on existing work rather than constituting the paper’s main new contribution. Nevertheless, establishing them correctly is essential because they provide the empirical inputs to the design-based decomposition that follows. The causal mediation analysis uses these causal estimates to separate the total income effect into the component operating through completed education years and the component remaining outside that channel.

3.1 Mendelian Independent Inheritance

I use variation away from parental values to estimate causal effects of having a higher versus lower Ed PGI. This approach leverages the fundamental randomness in genetic inheritance, first described by Mendel’s laws ([Mendel 1865](#)).

Each parent has 0, 1, or 2 copies of each genetic variant — corresponding to presence on each of their chromosomes. When a father and mother’s sex cells are formed, exactly one of the two copies is sorted into each reproductive cell, and is then inherited by their resulting child. A parent with one copy therefore passes it on half the time, whereas a parent with 0 or 2 copies transmits their genetic variant with certainty. At conception, the child receives one draw from the mother and one from the father, so the child’s copy-count is simply the sum of two independent draws.

Write $x_{i,j}^{\text{Child}}$ for the number of copies of variant j that individual i has (0, 1, or 2), and $x_{m(i),j}^{\text{Mother}}, x_{f(i),j}^{\text{Father}}$ for their respective parents $m(i), f(i)$. If the mother has two copies of variant j then they bequeath one copy to their child with certainty; if the father has one, then they bequeath one copy with probability $\frac{1}{2}$; if no copies, then they bequeath none with certainty. The child then receives their other copy from their father, by the same process, giving the following representation in expectation,

$$\mathbb{E} [x_{i,j}^{\text{Child}} \mid x_{m(i),j}^{\text{Mother}}, x_{f(i),j}^{\text{Father}}] = \frac{1}{2} (x_{m(i),j}^{\text{Mother}} + x_{f(i),j}^{\text{Father}}).$$

After conception, the child has inherited genetic variants from both parents — one of the 4 possible discrete realisations, not the expected value. This result came about purely from

random chance, conditional on at least one of parent having exactly one copy of the SNP, $x_{m(i),j}^{\text{Mother}} = 1$ and/or $x_{f(i),j}^{\text{Father}} = 1$.⁶

3.2 Mendelian Independent Inheritance in the Ed PGI

Moving from a single variant to the PGI simply means summing across the same random-segregation for all the variants. The Ed PGI is constructed by summing the 0, 1, or 2 copy indicators across the $J \approx 8,000$ different SNPs, each multiplied by the externally estimated weights w_j that capture how strongly each SNP predicts education years.

$$Z_i^{\text{Child}} = \sum_{j=1}^J w_j x_{i,j}^{\text{Child}},$$

and similarly for child i 's parents, to give $Z_{m(i)}^{\text{Mother}}$ and $Z_{f(i)}^{\text{Father}}$ from summing across binary values $x_{m(i),j}^{\text{Mother}}$ and $x_{f(i),j}^{\text{Father}}$, respectively.

Because every variant copy enters the child's genome through an independent random draw (conditional on parental values), the child's Ed PGI is the weighted sum of thousands of such draws. This means the Ed PGI can be decomposed into the component that was inherited as a result of a random draws from parent(s) who had 1 copy of the variant, and the certainty component from draws where parent(s) had 0 or 2 copies.

$$Z_i^{\text{Child}} = \mathbb{E} [Z_i^{\text{Child}} \mid Z_{m(i)}^{\text{Mother}}, Z_{f(i)}^{\text{Father}}] + u_i, \quad (1)$$

where u_i is a mean zero random error term arising from Mendelian independent assortment. The only determining factor for u_i is the randomness of inheriting exact counts of SNPs $x_{i,j}^{\text{Child}}$ from their parents, relative to what they were expected to inherit, $\frac{1}{2}(x_{m(i),j}^{\text{Mother}} + x_{f(i),j}^{\text{Father}})$.

Ideally, the expected inheritance is constructed from observed parental Ed PGI values, $Z_{m(i)}^{\text{Mother}}, Z_{f(i)}^{\text{Father}}$. However, parental PGI values are not observed for a large enough sample in the UKB for statistical inference. Thus, I follow [Young et al. \(2022\)](#) and leverage siblings' genetic data to impute parental PGI values. The expected Ed PGI inheritance, therefore, is calculated using the same within-family variation that early genetic studies did by examining within family differences. An individual's Ed PGI is drawn from the same distribution as their sibling, because they had the same parents; in absence of observing both parents directly, this approach uses the sibling to infer this distribution.


⁶Note both parents having 0 or 2 copies of variant j' leaves $x_{i',j'}^{\text{Child}} = \frac{1}{2}(x_{m(i'),j'}^{\text{Mother}} + x_{f(i'),j'}^{\text{Father}})$, so that individual i' inherited no random variation at variant j' .

This approach is an acceptable substitute to the ideal case of using observed parents' data, because of the simple fact that children and parents are two independent draws from the same parental genetic distribution. Write $Z_{s(i)}^{\text{Sibling}}$ for i 's sibling's Ed PGI value, and $Z_{i,s(i)}^{\text{Parental}}$ for the imputed parental value, calculated using genetic data from the individual and their observed sibling. In expectation, a child and their observed sibling's Ed PGI together proxy for their unobserved parental PGI values.

$$\mathbb{E} [Z_{i,s(i)}^{\text{Parental}}] = \frac{1}{2} \left(Z_{m(i)}^{\text{Mother}} + Z_{f(i)}^{\text{Father}} \right) = \mathbb{E} [Z_i^{\text{Child}} \mid Z_{m(i)}^{\text{Mother}}, Z_{f(i)}^{\text{Father}}].$$

However, because siblings collectively reveal approximately 75% of the parental genetic variation, the resulting imputed values contain measurement error compared to the ideal (but unobserved) parental mean. The sibling-imputed parental mean PGI can be viewed as a proxy for the true parental mean PGI, inflating the variance in error term u_i in Equation (1). This error reduces precision of using the imputed proxy measure, compared to true parental PGI values.

Figure 4: Distribution of the Ed PGI, Relative to Imputed Parental Decile.



Note: This figure shows the distribution of Ed PGI, Z_i^{Child} , and the Ed PGI minus imputed parental mean, $Z_i^{\text{Child}} - \mathbb{E} [Z_i^{\text{Child}} \mid Z_{m(i)}^{\text{Mother}}, Z_{f(i)}^{\text{Father}}]$, conditional on the decile of the parental imputed Ed PGI. Parental imputed Ed PGI decile 1 shows the distributions for the bottom 10% of parental imputed values, decile 2 shows distribution among 10–20%, and so on.

Children with a higher Ed PGI value inherit this from their parents, and so had parents with higher Ed PGI values. Figure 4 shows this tendency, where children with parents

in the top decile of the imputed Ed PGI distribution have a mean Ed PGI 1.5 standard deviations above the mean — and similarly for those at the bottom.⁷ This means parental Ed PGI values, with their effects on the children and the environment they are raised in, are a meaningful confounder for the effects of the child’s own Ed PGI on education and labour market outcomes. Panel B shows that, after conditioning for parental imputed values, there is no correlation across the distribution; this reflects the fact that Mendelian independent assortment induces random quasi-random variation in the child’s realised Ed PGI.

3.3 Causal Identification Strategy

Let $D_i(z)$ denote the potential education level and $Y_i(z, D_i(z))$ labour market outcome for individual i , under the hypothetical alternative realisation that i had Ed PGI value z . Observed education and wages/income data are equal to the potential outcomes for their realised Ed PGI value, Z_i^{Child} .

$$D_i = D_i(Z_i^{\text{Child}}), \quad Y_i = Y_i(Z_i^{\text{Child}}, D_i(Z_i^{\text{Child}})).$$

Potential labour market outcomes $Y_i(z, D_i(z))$ are a function of both the Ed PGI value, and resulting education $D_i(z)$. This distinction is not important for measuring average genetic effects in this section, but will become the main focus in [Section 5](#).

Target Parameters. This study aims to estimate the average causal effect of having a higher Ed PGI, which is summarised by the following marginal average effects — averaged across the distribution of Z_i^{Child} :⁸


$$\text{Average education effect} := \mathbb{E} \left[\frac{\partial}{\partial z} D_i(z) \right], \quad \text{Average total effect} := \mathbb{E} \left[\frac{\partial}{\partial z} Y_i(z, D_i(z)) \right].$$

The average education effect is the average marginal change in education induced by a marginal increase in the Ed PGI; The average total effect is the corresponding marginal change in labour-market income, allowing education to respond to the Ed PGI, *ceteris paribus*.⁹

⁷Panel A of [Figure 4](#) replicates the tendency shown in [Houmark et al. \(2024, Figure 1\)](#), with a new data source.

⁸Throughout, I assume that both $D_i(z)$ and $Y_i(z, D_i(z))$ are differentiable for any z across the distribution of Z_i^{Child} , together with standard regularity conditions permitting differentiation under the expectation.

⁹Under the linear specification used below, these marginal effects are constant. Because the Ed PGI is standardised, the regression coefficients can therefore be interpreted as the effects of a one-standard-deviation increase in the Ed PGI.

Figure 5: Structural Causal Model for Ed PGI Effects.

Note: This figure shows the structural causal model, representing the two causal effects estimated in this section. The causal effects of the child’s Ed PGI on education and later labour market outcomes are identified by controlling for parental (imputed) Ed PGI values.

Identifying Assumptions. I rely on controlling for parental (imputed) Ed PGI values, for causal inference. This yields estimates of causal effects if a child’s Ed PGI is quasi-randomly assigned, conditional on their parents’ Ed PGI mean — the child’s expected inheritance. For any z in the range of the Ed PGI distribution,

$$Z_i^{\text{Child}} \perp\!\!\!\perp D_i(z), Y_i(z, D_i(z)) \mid Z_{i,s(i)}^{\text{Parental}}. \quad (2)$$

This identifying assumption is credible precisely because of the Mendelian independent variation described above. The deviation away from expected inheritance represents the difference between a child’s realised Ed PGI and the Ed PGI they were expected to inherit given their imputed parents’ PGI values. This deviation arises solely from the random segregation of chromosomes during meiosis: at each of the thousands of SNPs entering the Ed PGI, whether a parent with one copy passes it on is determined by an independent coin flip, with no connection to the family environment, the child’s later circumstances, or any other factor that might influence education or labour market outcomes. A child who drew a higher Ed PGI than their expected inheritance is therefore not systematically different in family background, parental environment, or any other confounding factor from one who drew a lower value; the only difference is the outcome of independent biological coin flips at conception.

Estimation Strategy. I estimate the average effects of Ed PGI values on education and labour market outcomes with the following regression equations.

$$D_i = \phi_1 + \pi Z_i^{\text{Child}} + \zeta_1 Z_{i,s(i)}^{\text{Parental}} + \boldsymbol{\varphi}'_1 \mathbf{X}_i + \eta_i, \quad (3)$$

$$Y_i = \phi_2 + \theta Z_i^{\text{Child}} + \zeta_2 Z_{i,s(i)}^{\text{Parental}} + \boldsymbol{\varphi}'_2 \mathbf{X}_i + \epsilon_i, \quad (4)$$

where π is the effect of the Ed PGI on education variables (average education effect), θ the effect on later labour market outcomes (average total effect). Inclusion of parental imputed Ed PGI values yields causal inference, and $\mathbf{X}_i$ includes a birth set of control variables including gender, sibling count, in city indicator, birth year and UKB data collection year fixed effects, and inclusion of other PGI values as controls.¹⁰ Including these other PGI controls isolates the effects of the Ed PGI minus contribution from SNPs highly associated with other outcomes.

Using standard linear models, the marginal effects are assumed constant across the Ed PGI distribution, so that Ordinary Least Squares (OLS) consistently estimates a single parameter coinciding with the average total effect. Without linearity, OLS recovers a variance-weighted average of marginal effects, placing disproportionate weight on individuals whose Ed PGI deviates most from its conditional mean; the linearity assumption is therefore what gives the estimates their interpretation as population averages.

Standard errors are clustered at the family unit level throughout. The UKB recruited participants at the household level, which means that adult siblings are disproportionately represented among those who share a family unit — either because they were co-resident at the time of recruitment or because the household-based design increased the probability that multiple siblings enrolled. This induces within-family correlation in the residuals that, if ignored, would cause conventional standard errors to understate sampling uncertainty. Clustering at the family unit level accounts for arbitrary within-cluster correlation, yielding valid inference under the UKB sampling design.

Limitations. Several limitations apply to the estimates presented here. First, the OLS estimates of average education effect and average total effects (Equation (3) and (4)) are not simple averages across individuals, but rather convex weighted averages of individual-level causal effects, with weights determined by the OLS projection (Angrist & Krueger 1999). Without assuming homogeneity of effects across the Ed PGI distribution, these estimates need not coincide with the population average total effect, and may place disproportionate

¹⁰The other PGI variables are shown in Summary Table 1, including estimated PGIs for ADHD, asthma, bipolar disorder, body mass index, diabetes (type 2), height, and schizophrenia.

weight on individuals whose Ed PGI deviates most from its conditional mean.

Second, the analysis sample is restricted to UKB participants of European ancestry who have at least one observed sibling, as described in [Section 2](#). Internal validity is sound within this sample, although the estimated average effects may not generalise to the broader British population, whose observed and unobserved structure may systematically differ from this selected group. Third, the sibling-based imputation of parental PGI values introduces additional measurement error relative to the ideal case of directly observing true parental genotypes, attenuating the estimated coefficients. This attenuation is compounded by the fact that the Ed PGI itself is constructed from GWAS weights estimated with error, further understating the true causal effects of genetic predisposition toward educational attainment.¹¹

Table 2: Estimates of Education and Total Genetic Effects.

Outcome:	Education Outcomes			Labour Market Outcomes		
	Education years	Age left school	University completion	Occupation hourly wage	Occupation annual income	Household income (midpoint imputed)
	(1)	(2)	(3)	(4)	(5)	(6)
Ed PGI	0.546 (0.035)	0.163 (0.036)	0.072 (0.005)	0.049 (0.004)	0.047 (0.008)	0.058 (0.008)
Parental Ed PGI	0.870 (0.049)	0.337 (0.048)	0.105 (0.007)	0.067 (0.006)	0.066 (0.010)	0.094 (0.011)
Outcome mean	13.835	16.670	0.334	2.839	3.064	3.989
Adjusted R^2	0.146	0.043	0.107	0.138	0.205	0.136
Observations	24,702	16,726	24,702	24,702	20,560	22,181

Note: This table shows estimates of the average causal effect of the Ed PGI on education variables (average education effect) and labour market outcomes (average total effect) — specification in [Equation \(3\)](#) and [\(4\)](#). The coefficient on the Ed PGI represents the average causal effect, and coefficient on the parental Ed PGI is a control parameter. Specification for estimates in this table including gender, sibling count, in city indicator, birth year and UKB data collection year fixed effects, and inclusion of other PGI values as controls. Appendix [Table 3](#) shows estimates varying the control sets, and with sibling fixed effects.

4 Education and Total Genetic Effects


The Ed PGI has substantial causal effects on both educational attainment and labour market outcomes. [Table 2](#) reports estimates of the average causal effects, identified by the quasi-random variation in genetic inheritance described in [Section 3](#). Each coefficient on the Ed PGI represents the average causal effect of a one standard deviation increase in the Ed PGI, conditional on the imputed parental Ed PGI and the full set of controls. The parental Ed PGI enters as a control throughout: it is correlated with both the child’s Ed PGI and the outcomes of interest, but does not represent a causal effect, as it absorbs the component of a child’s Ed PGI that reflects inherited family environment rather than random genetic variation.

The Ed PGI has a meaningful causal effect on education, with the largest effects operating through formal qualification thresholds. A one standard deviation increase in the Ed PGI causes, on average, an additional 0.55 years of education (column 1; SE 0.02), against a sample mean of 13.84 years, and raises the probability of completing university by 7 percentage points on a baseline completion rate of 33 percent — a relative increase of around one fifth (column 3). Age of leaving full-time education, a common measure of education in the UK given that legal school-leaving age is defined by age rather than qualification, is observed for a subset of 16,728 participants; the estimated effect of 0.16 years (column 2; SE 0.02) is smaller than the qualification-based measure, reflecting that the Ed PGI operates more strongly through the decision to pursue higher qualifications than through the legally-determined school-leaving margin.

The Ed PGI has a consistent and precisely estimated causal effect on earnings across all three income measures. A one standard deviation increase in the Ed PGI raises hourly wages by approximately 5 percent (column 4) and annual occupation income by 5 percent (column 5), with a slightly larger effect of 6 percent on household income (column 6). These estimates capture the total average causal effect of the Ed PGI on labour market outcomes — combining any direct biological pathway with effects mediated through the additional education documented above.

The main estimates are robust to alternative specifications, and the main causal estimates are half as large as raw correlation between the Ed PGI for all outcomes under study. [Table 3](#) reports estimates under alternative specifications, and this raw correlation. Panel A shows the

¹¹The Ed PGI weights used in this paper (from [Okbay et al. 2022](#)) are calculated with measurement error, leading to attenuation bias relative to an infeasible truth Ed PGI with oracle weights. See [van Kippersluis et al. \(2023\)](#) for further discussion on measurement error of this kind.

Figure 6: Effect of Ed PGI on Education and Wages.

Note: This figure shows the correlation between Ed PGI and education years, as well as occupation coded hourly wages. The shifted lines are the estimates of the causal effect, estimated by controlling for parental imputed Ed PGI values.

raw correlations between the Ed PGI and outcomes, without conditioning on parental PGI or any controls; these are roughly twice the size of the causal estimates, illustrating the extent of confounding when parental genetics are not accounted for. Panel B re-estimates the main specification without the additional control variables, and finds estimates nearly identical to those in [Table 2](#), confirming that the causal identification rests on the parental PGI rather than the auxiliary controls. Including the other PGI for other outcomes as controls make little difference to the causal estimates, which aligns with specifications reporting correlations with the Ed PGI that do attenuate with inclusion of other PGIs (e.g., [Bolyard & Savelyev 2025](#)). Panel C replaces the parental PGI with sibling fixed effects, which partial out all shared family characteristics directly; the resulting estimates are again very close to the main results across education and labour market outcomes, lending further support to the identifying assumption that within-family variation in the Ed PGI is quasi-randomly assigned.

These findings closely replicate the main empirical specification and results in [Carvalho \(2025\)](#). The Ed PGI has meaningful causal effects on education and labour-market outcomes, while conditioning on the imputed parental Ed PGI substantially reduces the raw associations. The purpose of this replication is to establish the two inputs required for the analysis that

Table 3: Different Specifications for Estimates of Average Education and Total Effects.

Outcome:	Education Outcomes			Labour Market Outcomes		
	Education years	Age left school	University completion	Occupation hourly wage	Occupation annual income	Household income
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Raw Correlation						
Ed PGI	1.103 (0.020)	0.361 (0.020)	0.140 (0.003)	0.092 (0.003)	0.081 (0.005)	0.109 (0.005)
Panel B: Including Parental Ed PGI with no Control Variables						
Ed PGI	0.565 (0.035)	0.161 (0.036)	0.074 (0.005)	0.050 (0.005)	0.042 (0.009)	0.056 (0.008)
Parental Ed PGI	0.855 (0.049)	0.327 (0.049)	0.105 (0.007)	0.067 (0.006)	0.062 (0.012)	0.084 (0.012)
Panel C: Including Family Fixed Effects						
Ed PGI	0.560 (0.039)	0.294 (0.053)	0.070 (0.006)	0.052 (0.005)	0.061 (0.010)	0.065 (0.010)
Outcome mean	13.835	16.670	0.334	2.839	3.064	3.989
Observations	24,702	16,726	24,702	24,702	20,560	22,181

Note: This table shows estimates of the average causal effect of the Ed PGI on education variables (average first-stage) and labour market outcomes (average partial effect) — specification in Equation (3) and (4). The coefficient on the Ed PGI represents the average causal effect, and coefficient on the parental Ed PGI is a control parameter. Specification s vary for each Panel, as described.

follows: an Ed PGI effect of approximately 0.55 years on education, and a total income effect of approximately 5 percent. Having established these inputs, I now turn to the paper’s main contribution: decomposing the total effect of the Ed PGI on earnings into the share operating through education and a residual direct effect.

5 Direct and Indirect Effects

The average total effect of the Ed PGI on labour market outcomes is the sum of two conceptually distinct channels. The first is a direct channel: genetic predisposition toward education may influence earnings independently of education itself, through traits such as cognitive ability, non-cognitive skills, or personality characteristics that are correlated with the Ed

PGI and valued by employers. The second is an indirect channel: a higher Ed PGI induces more education, and that additional education generates a labour market return.

The size and robustness of the labour market effects documented in [Section 4](#) has led to a common interpretation in the literature that at least part of the Ed PGI’s influence must operate through the direct channel. This view is often motivated by the fact that the Ed PGI predicts scores on tests aiming to measure fluid intelligence ([Carvalho 2025](#)), by results showing earnings associations that persist even after conditioning on education ([Papageorge & Thom 2020](#), [Bryson et al. 2025](#)), and in work arguing that the Ed PGI captures skills related to financial literacy and other non-educational attributes ([Barth et al. 2020](#)). Separating these channels matters both for understanding the biological and social mechanisms behind genetic effects on earnings.

This section uses a Causal Mediation (CM) framework to decompose the average total effect of the Ed PGI into its direct and indirect components, where education years serves as the mediating variable. The key empirical challenge is that causal mediation analysis requires not only that the Ed PGI is quasi-randomly assigned (conditional on parental values), but also that the mediator is quasi-randomly assigned conditional on treatment, which is not credible for education in this setting.¹² UKB participants were not randomly assigned education, so it takes further steps to identify the returns to an additional year of education. Based on this insight, I combine credible estimates of the average effect of the Ed PGI on education together with estimates of the returns to education to identify the direct and indirect effect channels.

5.1 Causal Mediation (CM) Framework

Target Parameters. The total average total effect of the Ed PGI on labour market outcomes can be decomposed into a direct and indirect component. Formally, the direct and indirect effects are defined as:

$$\begin{aligned} \text{Direct Effect} &= \mathbb{E} \left[\frac{\partial}{\partial z} Y_i(z, D_i(Z_i^{\text{Child}})) \right], \\ \text{Indirect Effect} &= \mathbb{E} \left[\frac{\partial}{\partial z} Y_i(Z_i^{\text{Child}}, D_i(z)) \right]. \end{aligned}$$

The direct effect captures the effect of the Ed PGI on earnings holding the education channel fixed; the indirect effect captures the effect that operates exclusively through the induced


¹²[Appendix A.5](#) expands on this point, connecting to the instrumental variables literature.

change in education. These sum to the average total effect θ estimated in [Section 4](#),

$$\underbrace{\frac{\partial}{\partial z} \mathbb{E} [Y_i(z, D_i(z))]}_{\text{Average total effect, } \theta} = \underbrace{\frac{\partial}{\partial z} \mathbb{E} [Y_i(z, D_i(Z_i^{\text{Child}}))]}_{\text{Direct effect}} + \underbrace{\frac{\partial}{\partial z} \mathbb{E} [Y_i(Z_i^{\text{Child}}, D_i(z))]}_{\text{Indirect effect}},$$

evaluated at z — again, a linearity assumption is assuming the average effect does not vary across the Z_i^{Child} distribution. [Figure 7](#) illustrates this decomposition in the structural causal model: the direct effect corresponds to the Ed PGI arrow bypassing education, and the indirect effect to the path running through education years and its labour market return.

Figure 7: Structural Causal Model for Ed PGI Effects.



Note: This figure shows the structural causal model behind direct and indirect effects, where $Z_i \rightarrow D_i$ represents a causal relationship of Z_i on D_i . The dotted line between education and income & wages represents the fact that education is not randomly assigned, and a strategy for identifying returns to education is needed for valid CM estimates.

Identification. Non-parametric CM identification requires that the mediator is quasi-randomly assigned conditional on treatment (see [Imai et al. 2010](#)). This assumption is not credible here: education years are shaped by family background, ability, and a host of other factors that independently influence earnings, so that the education–earnings relationship might not represent the average causal effect for mediator compliers in the UKB sample (even after conditioning on the Ed PGI). However, it is possible to parametrically identify the direct and indirect effects with a valid estimate for the causal returns to education, by the chain rule.¹³

$$\underbrace{\frac{\partial}{\partial z} \mathbb{E} [Y_i(Z_i^{\text{Child}}, D_i(z))]}_{\text{Indirect Effect}} = \underbrace{\frac{\partial}{\partial z} \mathbb{E} [D_i(z)]}_{\text{Average education effect, } \pi} \times \underbrace{\frac{\partial}{\partial d} \mathbb{E} [Y_i(Z_i^{\text{Child}}, d)]}_{\text{Returns to education, } \beta}.$$

¹³This additionally requires the smoothness and regularity conditions assumed in [Section 3](#).

The indirect effect equals the average education effect π — credibly identified by the Mendelian design — multiplied by the returns to an additional year of education, β . With heterogeneous treatment effects, β represents the average education returns for people pushed into more years of education.¹⁴ The direct effect then follows as average total effect minus the indirect effect,

$$\text{Direct Effect} = \theta - \pi\beta.$$

This parametric identification also requires that no individuals have a negative education effect (i.e., a monotonicity condition), avoiding a negative weighting issue in the mediation indirect effects (Blackwell et al. 2024).

Estimation. I estimate the mediation second-stage from the following specification:

$$Y_i = \alpha + \beta D_i + \gamma Z_i^{\text{Child}} + \delta Z_i^{\text{Child}} D_i + \zeta Z_{i,s(i)}^{\text{Parental}} + \boldsymbol{\varphi}' \mathbf{X}_i + \varepsilon_i, \quad (5)$$

where β captures average returns to education, γ the residual direct effect of the Ed PGI on earnings conditional on education, and δ the interaction between the Ed PGI and education years. Since Z_i^{Child} is normalised to have mean zero, the interaction term $\delta Z_i D_i$ has mean zero across the sample, so that β recovers the average returns to education evaluated at the population mean of the Ed PGI.

Standard errors for the mediation estimates are obtained via a cluster bootstrap, resampling family units rather than individual observations. This preserves the within-family correlation in the residuals arising from the UKB sampling design, while also accommodating the non-linearity of the indirect effect estimator — which, as the product of two estimated quantities, does not have a closed-form standard error under heteroscedasticity. Each bootstrap draw resamples family units with replacement, estimates the mediation first- and second-stage parameters, and computes the implied direct and indirect effects; standard errors are the standard deviations of the resulting bootstrap distributions.

Limitations. Several limitations apply to the CM estimates presented in this section. First, the OLS estimates of the average total effect and average education effect are variance-weighted averages of individual-level effects under heterogeneity (as noted in Section 3). Without the linearity assumption, these average effects need not coincide with population

¹⁴Note that this is a parametrically identified model, because the correlational measure (and accompanying sensitivity analysis) is meant to inform the local average effect for $D_i(Z_i^{\text{Child}})$ compliers — but does not non-parametrically identify it.

average effects. Second, and more substantively, the indirect effect in [Section 5.1](#) implicitly assumes that the returns to education are homogeneous: individuals induced into additional schooling by a higher Ed PGI face the same average returns as the broader population. In the language of CM, this is an assumption about the education compliers — those whose schooling decisions are shifted by the Ed PGI — whose returns may differ from the population average if, for example, education-linked genes select individuals into different schooling margins with different labour market rewards. The validity of the CM estimates rests on the credibility of β as an estimate of the returns to education relevant for this complier population. The only remaining step for identifying CM effects is to provide credible estimates of the causal returns to education in the UKB sample.

5.2 Returns to Education

Correlational estimates of education returns are the most natural starting point. A large literature finds that OLS estimates of the return to an additional year of schooling are remarkably stable; they are typically estimated in the range of 5–10 percent in log earnings, and close to estimates from quasi-experimental designs ([Angrist & Krueger 1991](#), [Harmon & Walker 1995](#), [Oreopoulos 2006](#)). Two forces likely explain this convergence: ability bias pushes OLS upward, while measurement error in reported schooling pushes it downward, and the two roughly cancel. Quasi-experimental designs, meanwhile, often identify effects among specific groups of marginal compliers, whose average education returns need not equal the population average ([Buscha & Dickson 2015](#)). The fact that OLS and quasi-experimental estimates land in a similar range is informative, not coincidental.

The correlational estimate for the labour market return to an additional year of education in the UKB analysis sample sits consistently at around 6 percent in log earnings across all three income measures (shown in [Table 4](#)). This estimate is stable across specifications: adding controls for the Ed PGI and its interaction with education years barely moves the point estimates, and including the parental Ed PGI changes nothing. The UKB correlational estimates are, in short, consistent with the broader literature.

Canonical labour economics work was concerned that the correlation between years of completed education and later life income would be confounded by unobserved ability ([Griliches 1977](#)). One may expect that genetic variants, contained in the Ed PGI, would be correlated with unobserved ability, and thus controlling for the Ed PGI may lead to lower correlational estimates for the returns to education. This is not the case in UKB data, and is consistent with follow-up research that uses casual methods such as instrumental variables and finds

Table 4: Returns to Education, Correlational Estimates.

Outcome:	Occupation hourly wage			Occupation annual income			Household income (midpoint imputed)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Education years	0.061 (0.001)	0.059 (0.001)	0.059 (0.001)	0.061 (0.001)	0.059 (0.001)	0.059 (0.001)	0.068 (0.002)	0.063 (0.002)	0.063 (0.002)
Ed PGI		0.010 (0.011)	0.000 (0.010)		0.008 (0.021)	-0.001 (0.021)		0.008 (0.013)	-0.015 (0.010)
Education years $\times$ Ed PGI		0.001 (0.001)	0.001 (0.001)		0.001 (0.001)	0.001 (0.001)		0.003 (0.001)	0.003 (0.001)
Parental Ed PGI			0.016 (0.003)			0.016 (0.005)			0.040 (0.007)
Collected Education Returns	0.061 (0.001)	0.059 (0.001)	0.059 (0.001)	0.061 (0.001)	0.059 (0.001)	0.059 (0.001)	0.068 (0.002)	0.064 (0.002)	0.063 (0.002)
R^2	0.328	0.332	0.332	0.273	0.274	0.274	0.204	0.208	0.208
Observation count	24,702	24,702	24,702	20,560	20,560	20,560	22,181	22,181	22,181

Note: This table shows simple OLS estimates for returns to education in the UK Biobank sibling sample — the relevant sample for causal genetic effects. All outcomes are in a log specification, and each column is a regression including the controls previously mentioned, with specification in [Equation \(5\)](#), varying inclusion of the Ed PGI and parental Ed PGI as indicated. Standard errors are clustered at the family level. Observation counts vary because each outcome has a different number of observations which are non-missing. The final row for education returns collects the partial effect of education years, including the interaction terms between education years and the Ed PGI. Point estimates are relatively similar to those using the MLSA rise as an instrument (see [Appendix A.4](#)), with much more precision.

estimates for returns to education similar as the correlational evidence (see [Card 1999](#) for an overview). The stability of the return to education after conditioning on the Ed PGI and parental Ed PGI therefore suggests that ability bias, at least insofar as it operates through genetic endowment, is not a material source of upward distortion in the correlational estimate.


Nonetheless, it would still not be credible to assume education years are randomly assigned, so that no correlational estimate is entirely credible on its own. The UKB does contain a plausible quasi-experiment: the 1972 rise in the British Minimum School Leaving Age (MLSA), which compelled children born after September 1957 to stay in school until age 16. While this sample lines up well with the UKB population, the resulting fuzzy RDD estimates of education returns (in the range of 4–6 percent in log earnings) carry standard errors too large to rule out either zero or substantially larger effects, and are therefore too imprecise

for use in the mediation analysis.¹⁵ Given this, the correlational estimates anchored in the broader UK literature offer a more useful basis for the parametric mediation analysis, and the sensitivity analysis below makes the dependence on any particular return value transparent.

5.3 Sensitivity Analysis

The indirect effect is the product of two quantities: the effect on education years and the return to an additional year of education. Since these multiply together, any uncertainty in the education return translates directly and proportionally into uncertainty about the share of the total genetic effect that operates through education. Rather than conditioning the mediation analysis on a single return value, I vary it across its plausible range. In addition, this range is defined not arbitrarily, but based on the previous labour economics on education returns.

Figure 8: Distribution of Economics Literature for Income Returns to Education.



Note: These figures show the estimated kernel density for the distribution of estimates for later-life income returns to education. Panel A refers to the implied distribution from 17 estimates using IVs in the UK (shown in full in Appendix Table A3), and Panel B for 182 estimates from 144 studies across the world. These estimates were collected by Patrinos & Psacharopoulos (2026).

The return to an extra year of education in the UK has been estimated many times in the economics literature, giving a rich empirical base for calibrating the sensitivity analysis. Figure 8 shows the distribution of 17 quasi-experimental UK estimates collected by Patrinos & Psacharopoulos (2026), listed in full in Appendix Table A3. The distribution is right-skewed:

¹⁵ Appendix A.4 details estimates for labour market education returns using the MLSA rise as a regression discontinuity, as previously considered by Barcellos et al. (2021).

most estimates fall between 4 and 8 percent, consistent with the correlational estimates in [Table 4](#), but a tail of larger estimates reaches to around 15 percent. These higher values tend to come from studies exploiting the 1947 compulsory schooling reform, which identifies returns among a different and older complier population than the UKB cohort. The worldwide distribution from 182 estimates across 144 studies, shown alongside, is broadly consistent with the UK evidence.

The first step of the sensitivity analysis asks a deliberately agnostic question: how do the mediation estimates change as the assumed education return varies across the full span of values that appear in the literature, from 0 to 20 percent? For each value in this range, the indirect effect follows directly as the product of the estimated genetic education effect and the assumed return, with the direct effect as the remainder. This step takes no stand on which return value is correct, and shows whether the mediation conclusions are stable or fragile across the plausible range.

The second step uses the empirical UK distribution more directly. I bootstrap across the 17 IV estimates in the economic literature (shown in Appendix [Table A3](#)), drawing a return value at each replication and computing the implied mediation decomposition. Repeating this across many draws produces a distribution of indirect and direct effect estimates that reflects both sampling uncertainty in the effect of the Ed PGI on education years and the genuine disagreement in the literature about the magnitude of education returns. Reporting the point estimate, and bootstrapped standard errors of this distribution gives a meta-analytic sensitivity bound anchored in what the literature actually finds, rather than treating all values across the plausible range as equally likely.

Together, the two steps answer different questions. The first shows the full range of possible conclusions. The second shows where the weight of evidence places them.

6 Direct and Indirect Effect Estimates

Education years account for the majority of the Ed PGI's effect on labour market earnings. At correlational returns to education of around 6 percent, roughly 65 to 75 percent of the total genetic effect on earnings operates through the education channel, with a modest residual direct effect. A sensitivity analysis varying the assumed return across the full range of credible estimates from the economics literature shows this conclusion is robust: for return values above 7 to 8 percent. These values are well within the mass of UK quasi-experimental estimates, and at these values the direct genetic effect becomes indistinguishable from zero,

with the education channel accounts for essentially the entire Ed PGI effect.

Table 5: Direct and Indirect Effect Estimates, using Correlational Education Returns.


Outcome:	Occupation hourly wage		Occupation annual income		Household income (midpoint imputed)	
	(1)	(2)	(3)	(4)	(5)	(6)
Education effect, π	0.565 (0.037)	0.551 (0.037)	0.572 (0.040)	0.557 (0.040)	0.569 (0.039)	0.559 (0.039)
Total genetic effect, θ	0.050 (0.005)	0.049 (0.005)	0.042 (0.009)	0.048 (0.008)	0.056 (0.009)	0.059 (0.008)
Education returns, β	0.058 (0.001)	0.059 (0.001)	0.064 (0.001)	0.060 (0.001)	0.071 (0.001)	0.064 (0.001)
Direct, Ed PGI effect $\theta - \pi\beta$	0.018 (0.004)	0.017 (0.004)	0.006 (0.009)	0.015 (0.008)	0.015 (0.009)	0.023 (0.008)
Indirect, education effect $\pi\beta$	0.033 (0.002)	0.033 (0.002)	0.036 (0.003)	0.033 (0.003)	0.041 (0.003)	0.036 (0.003)
Proportion mediated through education, $\frac{\pi\beta}{\theta}$	0.651 (0.056)	0.653 (0.059)	0.853 (0.224)	0.693 (0.129)	0.725 (0.130)	0.611 (0.091)
Controls included?	No	Yes	No	Yes	No	Yes
Observation count	24,743	24,702	20,592	20,560	22,215	22,181

Note: This table shows the mediation point estimates, of the effect of the Ed PGI going through education years (indirect effect) and a direct genetic effect. Specification controls for parental Ed PGI and uses correlation returns to education, as mentioned in the previous section. All outcomes are in a log specification, and each column varies inclusion of the control variables as indicated, and standard errors on the mediation estimates are calculated by the bootstrap with 10,000 bootstrap samples. Observation counts vary because each outcome has a different number of observations which are non-missing.

Education mediates the majority of the Ed PGI's effect on earnings across all three income measures (Table 5). Using correlational returns to education, the indirect channel accounts for around 65 percent of the total genetic effect on hourly wages (column 2) and 69 percent on annual income (column 4), with a slightly wider range for household income (column 6). The residual direct genetic effect is small across all outcomes, around a third the size of the indirect effect, and the education channel dominates in each preferred specification. The no-controls annual income estimate is a noticeable outlier at 85 percent mediated, though its wide standard error warrants treating the controlled specifications as the preferred estimates.

The conclusion that education years dominate is robust across the full range of credible returns to education. Figure 9 traces the decomposition for log occupational hourly wages as

Figure 9: Sensitivity Analysis, Effect of Ed PGI on Log Occupational Wages via Education Years.



Note: This figure shows how the indirect (education) and direct effect estimates vary using difference values for average returns to education. The total effect refers to the average effect of the Ed PGI on log occupational hourly wages after controlling for parental Ed PGI values. The shaded regions refer to 95% confidence intervals, calculated with bootstrapped standard errors clustered at the sibling family unit. The dotted line labelled “Correlational Estimate” shows the position of the correlational returns to education estimate, presented in [Table 4](#).


the assumed return varies from zero to 20 percent.¹⁶ At returns above roughly 7 to 8 percent returns to education, the direct genetic effect becomes statistically indistinguishable from zero, and education years account for the entirety of the total genetic earnings effect. Above 10 percent the direct effect point estimate turns negative, which is difficult to reconcile with any plausible mechanism; average returns for the complier population induced by the Ed PGI are unlikely to reach that level. The lower bound is zero by construction, because if extra education years carry no labour market return, none of the genetic effect can operate through it. If average returns to education were negative, then the Ed PGI channel through education would also be negative — which is unlikely.

The same pattern holds across all three income measures. [Figure 10](#) plots the ratio of the indirect effect to the total genetic effect for each outcome as the assumed return varies, and the picture is consistent: for occupational hourly wages, annual income, and household income alike, education years account for the full total genetic effect at returns of around 7 to 8 percent. The three outcomes track closely together across the range, reinforcing that

¹⁶The average total effect is the same as that presented in [Section 4](#), and is not varied as part of this sensitivity analysis.

Figure 10: Sensitivity Analysis, Proportion of Ed PGI Effect on Income Outcomes via Education Years.

Indirect Effect / Total Effect



Note: This figure shows how the indirect (education) effect divided by total genetic effect, for each income measurement, showing percent of the Ed PGI effect mediated through education years. “Occ” refers to the word “occupational,” due to space constraints. The shaded regions refer to 95% confidence intervals, calculated with bootstrapped standard errors. Specifications involve controlling for parental Ed PGI values, and including the other PGI values as controls. The dotted line labelled “Correlational Estimate” shows the position of the correlational returns to education estimate, presented in [Table 4](#).

the dominance of the education years channel is not specific to one of the available measures of later-life labour market income in the UKB.

Sampling directly from the empirical distribution of UK quasi-experimental returns to education shifts the central mediation share further toward the education years channel [Table 6](#). Only one of the 17 IV estimates in the literature falls below the OLS anchor of around 6 percent, so the bootstrapped average return of 7.6 percent exceeds the correlational baseline; at this return, the direct genetic effect is indistinguishable from zero across all three income measures, and education years account for around 85 to 89 percent of the total genetic earnings effect. The confidence intervals on the mediation shares are wide, but this uncertainty is inherited almost entirely from the dispersion of return estimates in the literature rather than from sampling noise in the genetic education effect. What the literature agrees on is that returns to education are at or above 7 percent; above that threshold, the case for a meaningful direct genetic effect on earnings does not survive.

Table 6: Direct and Indirect Effect Estimates, using Bootstrap Distribution of Education Returns Estimates.

Outcome:	Occupation hourly wage	Occupation annual income	Household income (midpoint imputed)
	(1)	(2)	(3)
Education and Total Effect Estimates:			
Total genetic effect, θ	0.049 (0.005)	0.048 (0.008)	0.059 (0.008)
Education returns, β	0.076 (0.036)	0.076 (0.036)	0.076 (0.036)
Mediation Estimates:			
Direct, Ed PGI effect $\theta - \pi\beta$	0.008 (0.021)	0.006 (0.021)	0.017 (0.021)
Indirect, education effect $\pi\beta$	0.042 (0.020)	0.042 (0.020)	0.042 (0.020)
Proportion mediated through education $\frac{\pi\beta}{\theta}$	0.847 (0.414)	0.889 (0.459)	0.719 (0.361)
Observation count	24,743	20,592	22,215

Note: This table shows the mediation point estimates, of the effect of the Ed PGI going through education years (indirect effect) and a direct genetic effect. Data are bootstrapped sampled from the UKB analysis sample, and a from the distribution of causal income returns to education, provided by [Patrinos & Psacharopoulos \(2026\)](#) and shown in Appendix [Table A3](#), to capture statistical uncertainty from both education returns distribution and the UKB genetic estimates. Specification controls for parental Ed PGI and uses correlation returns to education, as mentioned in the previous section. All outcomes are in a log specification, and standard error estimates are calculated by the bootstrap with 1,000 bootstrap samples. Observation counts vary because each outcome has a different number of observations which are non-missing.

The Ed PGI’s effect on labour market earnings operates primarily through education years. At correlational returns to education of around 6 percent, roughly 65 to 75 percent of the total genetic effect is mediated through the education years channel, with a residual direct effect of 35 to 25 percent. Under the majority of plausible returns to education from the UK economics literature, this share rises further, and the direct genetic effect becomes undetectable in UKB data.

6.1 Discussion

These results partially replicate and extend [Carvalho \(2025\)](#)’s findings on the Ed PGI and labour market outcomes, and provide an account of the primary mechanism. The total genetic

effects on education and earnings are broadly consistent with that paper’s estimates; the mediation analysis goes further by showing that education decisions, which genetic inheritance shapes, are the primary route through which the Ed PGI translates into earnings inequality. Individuals with a higher Ed PGI attend more years of education, and it is these additional education years, valued by the labour market, that accounts for the majority of their earnings advantage. The direct genetic effect, whatever its biological basis, is at most a secondary force.

The education years mediating mechanism here is, if anything, a conservative account of the full role of education. The UKB records no measures of education attainment quality, such as examination scores, ranking of university attended, or field of study; the mediator of education years completed captures only the education quantity margin. If the Ed PGI also raises education attainment quality, then the education channel (combining quality and quantity) would be larger still, and the residual direct genetic effect correspondingly smaller.

A second limitation is that the analysis is restricted to UKB participants with at least one sibling in the data. This sample cannot be generalised to the broader British population, and the sibling-imputation procedure introduces measurement error relative to the ideal case of observing both parental Ed PGI values directly. A subtler issue is that when parents are far apart in the Ed PGI distribution, the range of values a child can inherit is wider; this additional variation in expected inheritance is not fully captured by the sibling-imputed procedure.

Future research would benefit from data that observes both parents’ Ed PGI values directly, removing the measurement error introduced by sibling imputation and improving the precision of the genetic education effect. Progress on the quality margin would require data measuring any of the education quality channels. Together, these extensions would sharpen both the causal estimates and the mediation analysis presented here.

7 Summary and Concluding Remarks

This paper uses quasi-random variation in genetic inheritance to estimate the causal effects of the Ed PGI on education and labour market outcomes, and decomposes those effects into a direct genetic channel and an indirect channel operating through education years in a causal mediation framework. The Ed PGI raises both years of education and later-life income by multiple measures, replicating the core findings of the recent literature ([Carvalho 2025](#)). A causal mediation analysis uses these findings as a basis, and at plausible returns to

education from the economics literature for Britain, the majority of the total Ed PGI income effect operates through the education years channel, with a residual direct genetic effect that is small and, under higher return values, indistinguishable from zero. The literature had noted suggestive evidence that direct effects may exist, but had not delivered a quantitative decomposition; this paper does so, and finds that the education years channel accounts for the bulk of the total effect across the plausible range of returns, with (at most) small direct effects.

References

- Angrist, J. D. & Krueger, A. B. (1991), ‘Does compulsory school attendance affect schooling and earnings?’, *The Quarterly Journal of Economics* **106**(4), 979–1014. <https://doi.org/10.2307/2937954>. 26
- Angrist, J. D. & Krueger, A. B. (1999), Chapter 23 – empirical strategies in labor economics, Vol. 3 of *Handbook of Labor Economics*, Elsevier, pp. 1277–1366. [https://doi.org/10.1016/S1573-4463\(99\)03004-7](https://doi.org/10.1016/S1573-4463(99)03004-7). 18
- Barcellos, S. H., Carvalho, L., Langa, K., Nimmagadda, S. & Turley, P. (2025), Education and dementia risk, Technical report, National Bureau of Economic Research. <https://doi.org/10.3386/w33430>. 44
- Barcellos, S. H., Carvalho, L. S. & Turley, P. (2023), ‘Distributional effects of education on health’, *Journal of Human Resources* **58**(4), 1273–1306. <https://doi.org/10.3368/jhr.59.2.0720-11064r1>. 44
- Barcellos, S. H., Carvalho, L. & Turley, P. (2021), The effect of education on the relationship between genetics, early-life disadvantages, and later-life ses, Technical report, National Bureau of Economic Research. <https://doi.org/10.3386/w28750>. 28, 44, 46
- Barth, D., Papageorge, N. W. & Thom, K. (2020), ‘Genetic endowments and wealth inequality’, *Journal of Political Economy* **128**(4), 1474–1522. <https://doi.org/10.1086/705415>. 23
- Barth, D., Papageorge, N. W., Thom, K. & Velásquez-Giraldo, M. (2022), Genetic endowments, income dynamics, and wealth accumulation over the lifecycle, Technical report, National Bureau of Economic Research. <https://doi.org/10.3386/w30350>. 3
- Beauchamp, J. P., Cesarini, D., Johannesson, M., van der Loos, M. J. H. M., Koellinger, P. D., Groenen, P. J. F., Fowler, J. H., Rosenquist, J. N., Thurik, A. R. & Christakis, N. A. (2011), ‘Molecular genetics and economics’, *Journal of Economic Perspectives* **25**(4), 57–82. <https://doi.org/10.1257/jep.25.4.57>. 2
- Biroli, P., Galama, T., von Hinke, S., van Kippersluis, H., Rietveld, C. A. & Thom, K. (2025), ‘The economics and econometrics of gene–environment interplay’, *The Review of Economic Studies* **93**(1), 144–180. <https://doi.org/10.1093/restud/rdaf034>. 3
- Blackwell, M., Ma, R. & Opacic, A. (2024), ‘Assumption smuggling in intermediate outcome tests of causal mechanisms’, *arXiv preprint arXiv:2407.07072*. <https://doi.org/10.48550/arXiv.2407.07072>. 3, 25
- Bolt, U., French, E., Warrier, V., Yang, Q. & Zhang, W. (2025), ‘Genetic endowments and lifetime earnings: Understanding the mechanisms’, *Working paper*. <https://www.econ.cam.ac.uk/sites/default/files/publication-cwpe-pdfs/cwpe2547.pdf>. 3

- Bolyard, A. & Savelyev, P. A. (2025), ‘Understanding the educational attainment polygenic index and its interactions with ses in determining health in young adulthood’, *Journal of Human Resources* . <https://doi.org/10.3368/jhr.0521-11685r3>. 21
- Bowden, J., Davey Smith, G. & Burgess, S. (2015), ‘Mendelian randomization with invalid instruments: effect estimation and bias detection through egger regression’, *International journal of epidemiology* **44**(2), 512–525. <https://doi.org/10.1093/ije/dyv080>. 49
- Bryson, A., Morris, T., Bann, D. & Wilkinson, D. (2025), ‘The gender wage gap across life: Effects of genetic predisposition towards higher educational attainment’, *Economics & Human Biology* **56**, 101471. <https://doi.org/10.1016/j.ehb.2025.101471>. 3, 23
- Buscha, F. & Dickson, M. (2015), The wage returns to education over the life-cycle: heterogeneity and the role of experience, Technical report, IZA Discussion Papers. 26
- Calonico, S., Cattaneo, M. D. & Titiunik, R. (2014), ‘Robust nonparametric confidence intervals for regression-discontinuity designs’, *Econometrica* **82**(6), 2295–2326. <https://doi.org/10.3982/ecta11757>. 48
- Calonico, S., Cattaneo, M. D. & Titiunik, R. (2015), ‘rdrbust: An r package for robust nonparametric inference in regression-discontinuity designs’, *The R Journal* **7**, 38–51. <https://doi.org/10.32614/RJ-2015-004>. 48
- Card, D. (1999), ‘The causal effect of education on earnings’, *Handbook of labor economics* **3**, 1801–1863. [https://doi.org/10.1016/s1573-4463\(99\)03011-4](https://doi.org/10.1016/s1573-4463(99)03011-4). 27
- Carvalho, L. S. (2025), ‘Genetics and socioeconomic status: Some preliminary evidence on mechanisms’, *Journal of Political Economy Microeconomics* **3**(3), 000–000. <https://doi.org/10.1086/732835>. 1, 2, 3, 6, 21, 23, 33, 34, 42
- Clark, D. & Royer, H. (2013), ‘The effect of education on adult mortality and health: Evidence from britain’, *American Economic Review* **103**(6), 2087–2120. <https://doi.org/10.1257/aer.103.6.2087>. 44, 46
- Conley, T. G., Hansen, C. B. & Rossi, P. E. (2012), ‘Plausibly exogenous’, *The Review of Economics and Statistics* **94**(1), 260–272. https://doi.org/10.1162/REST_a_00139. 49
- Cunha, F. & Heckman, J. (2007), ‘The technology of skill formation’, *American Economic Review* **97**(2), 31–47. <https://doi.org/10.1257/aer.97.2.31>. 2
- Davies, N. M., Dickson, M., Davey Smith, G., Van Den Berg, G. J. & Windmeijer, F. (2018), ‘The causal effects of education on health outcomes in the uk biobank’, *Nature human behaviour* **2**(2), 117–125. <https://doi.org/10.1038/s41562-017-0279-y>. 44
- Gelman, A. & Imbens, G. (2019), ‘Why high-order polynomials should not be used in regression discontinuity designs’, *Journal of Business & Economic Statistics* **37**(3), 447–456. <https://doi.org/10.1080/07350015.2017.1366909>. 48

- Green, D. P., Ha, S. E. & Bullock, J. G. (2010), ‘Enough already about “black box” experiments: Studying mediation is more difficult than most scholars suppose’, *The Annals of the American Academy of Political and Social Science* **628**(1), 200–208. <https://doi.org/10.1177/0002716209351526>. 3
- Griliches, Z. (1977), ‘Estimating the returns to schooling: Some econometric problems’, *Econometrica* **45**(1), 1–22. <https://doi.org/10.2307/1913285>. 26
- Harmon, C. & Walker, I. (1995), ‘Estimates of the economic return to schooling for the united kingdom’, *The American Economic Review* **85**(5), 1278–1286. <https://doi.org/10.1920/wp.ifs.1995.9512>. 26
- Hill, W. G., Goddard, M. E. & Visscher, P. M. (2008), ‘Data and theory point to mainly additive genetic variance for complex traits’, *PLoS genetics* **4**(2), e1000008. <https://doi.org/10.1371/journal.pgen.1000008>. 8
- Hogan-Hennessy, S. (2025), ‘Causal Mediation in Natural Experiments.’. <https://doi.org/10.48550/arXiv.2508.05449>. 3
- Houmark, M. A., Ronda, V. & Rosholm, M. (2024), ‘The nurture of nature and the nature of nurture: How genes and investments interact in the formation of skills’, *American Economic Review* **114**(2), 385–425. <https://doi.org/10.1257/aer.20220456>. 3, 16
- Huber, M. (2020), ‘Mediation analysis’, *Handbook of labor, human resources and population economics* pp. 1–38. https://doi.org/10.1007/978-3-319-57365-6_162-1. 3
- Imai, K., Keele, L. & Yamamoto, T. (2010), ‘Identification, inference and sensitivity analysis for causal mediation effects’, *Statistical Science* pp. 51–71. <https://doi.org/10.1214/10-sts321>. 3, 24
- Kolesár, M., Chetty, R., Friedman, J., Glaeser, E. & Imbens, G. W. (2015), ‘Identification and inference with many invalid instruments’, *Journal of Business & Economic Statistics* **33**(4), 474–484. <https://doi.org/10.1080/07350015.2014.978175>. 49
- Kweon, H., Burik, C. A., Ning, Y., Ahlskog, R., Xia, C., Abner, E., Bao, Y., Bhatta, L., Faquih, T. O., de Feijter, M. et al. (2025), ‘Associations between common genetic variants and income provide insights about the socio-economic health gradient’, *Nature human behaviour* pp. 1–12. <https://doi.org/10.1038/s41562-024-02080-7>. 6, 42, 46
- Manichaikul, A., Mychaleckyj, J. C., Rich, S. S., Daly, K., Sale, M. & Chen, W.-M. (2010), ‘Robust relationship inference in genome-wide association studies’, *Bioinformatics* **26**(22), 2867–2873. <https://doi.org/10.1093/bioinformatics/btq559>. 41
- Martin, A. R., Gignoux, C. R., Walters, R. K., Wojcik, G. L., Neale, B. M., Gravel, S., Daly, M. J., Bustamante, C. D. & Kenny, E. E. (2017), ‘Human demographic history impacts genetic risk prediction across diverse populations’, *The American Journal of Human Genetics* **100**(4), 635–649. <https://doi.org/10.1016/j.ajhg.2017.03.004>. 7

- Mendel, G. (1865), ‘Versuche uber pflanzen-hybriden’, *Vorgelegt in den Sitzungen* . <https://doi.org/10.5962/bhl.title.61004>. 13
- Millwood, I. Y., Im, P. K., Bennett, D., Hariri, P., Yang, L., Du, H., Kartsonaki, C., Lin, K., Yu, C., Chen, Y. et al. (2023), ‘Alcohol intake and cause-specific mortality: conventional and genetic evidence in a prospective cohort study of 512 000 adults in china’, *The Lancet Public Health* **8**(12), e956–e967. [https://doi.org/10.1016/s2468-2667\(23\)00217-7](https://doi.org/10.1016/s2468-2667(23)00217-7). 7
- Muslimova, D., Van Kippersluis, H., Rietveld, C. A., Von Hinke, S. & Meddens, S. F. W. (2025), ‘Gene-environment complementarity in educational attainment’, *Journal of Labor Economics (Forthcoming)* . <https://doi.org/10.1086/734087>. 3, 41, 42
- Okbay, A., Wu, Y., Wang, N., Jayashankar, H., Bennett, M., Nehzati, S. M., Sidorenko, J., Kweon, H., Goldman, G., Gjorgjieva, T. et al. (2022), ‘Polygenic prediction of educational attainment within and between families from genome-wide association analyses in 3 million individuals’, *Nature genetics* **54**(4), 437–449. <https://doi.org/10.3410/f.733675894.793549439>. 1, 4, 8, 20, 46
- Oreopoulos, P. (2006), ‘Estimating average and local average treatment effects of education when compulsory schooling laws really matter’, *American Economic Review* **96**(1), 152–175. <https://doi.org/10.1257/000282806776157641>. 26, 44
- Papageorge, N. W. & Thom, K. (2020), ‘Genes, education, and labor market outcomes: evidence from the health and retirement study’, *Journal of the European Economic Association* **18**(3), 1351–1399. <https://doi.org/10.1093/jeea/jvz072>. 3, 23
- Patrinos, H. A. & Psacharopoulos, G. (2026), ‘Causal returns to education’, *International Journal of Educational Development* **122**, 103565. <https://doi.org/10.1016/j.ijedudev.2026.103565>. 1, 2, 28, 33, 51
- R Core Team (2023), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>. 41
- Sanz-de Galdeano, A. & Terskaya, A. (2025), ‘Sibling differences in genetic propensity for education: How do parents react?’, *Review of Economics and Statistics* **107**(5), 1356–1370. https://doi.org/10.1162/rest_a_01326. 3
- Schork, A. J., Peterson, R. E., Dahl, A., Cai, N. & Kendler, K. S. (2022), ‘Indirect paths from genetics to education’, *Nature genetics* **54**(4), 372–373. <https://doi.org/10.1038/s41588-021-00999-5>. 3
- Smith, A. (1776), *The Wealth of Nations*, Vol. 1. 1
- Spiller, W., Slichter, D., Bowden, J. & Davey Smith, G. (2019), ‘Detecting and correcting for bias in mendelian randomization analyses using gene-by-environment interactions’, *International journal of epidemiology* **48**(3), 702–712. <https://doi.org/10.1093/ije/dyy204>. 49

- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M. et al. (2015), ‘Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age’, *PLoS medicine* **12**(3), e1001779. <https://doi.org/10.1371/journal.pmed.1001779>. 4
- Uffelmann, E., Huang, Q. Q., Munung, N. S., De Vries, J., Okada, Y., Martin, A. R., Martin, H. C., Lappalainen, T. & Posthuma, D. (2021), ‘Genome-wide association studies’, *Nature Reviews Methods Primers* **1**(1), 59. <https://doi.org/10.1038/s43586-021-00056-9>. 8
- van Kippersluis, H., Biroli, P., Dias Pereira, R., Galama, T. J., von Hinke, S., Meddens, S. F. W., Muslimova, D., Slob, E. A., de Vlaming, R. & Rietveld, C. A. (2023), ‘Overcoming attenuation bias in regressions using polygenic indices’, *Nature communications* **14**(1), 4473. <https://doi.org/10.1038/s41467-023-40069-4>. 20
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., Golemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T. L., Miller, E., Bache, S. M., Müller, K., Ooms, J., Robinson, D., Seidel, D. P., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K. & Yutani, H. (2019), ‘Welcome to the tidyverse’, *Journal of Open Source Software* **4**(43), 1686. <https://doi.org/10.21105/joss.01686>. 41
- Widding-Havneraas, T., Demange, P. A., Zachrisson, H. D., Borgen, N., Ystrom, E. & Elwert, F. (2026), ‘Estimating returns to education using the genetic lottery’, *Proceedings of the National Academy of Sciences* **123**(15). <https://doi.org/10.1073/pnas.2537049123>. 48
- Young, A. I., Nehzati, S. M., Benonisdottir, S., Okbay, A., Jayashankar, H., Lee, C., Cesarini, D., Benjamin, D. J., Turley, P. & Kong, A. (2022), ‘Mendelian imputation of parental genotypes improves estimates of direct genetic effects’, *Nature genetics* **54**(6), 897–905. <https://doi.org/10.1038/s41588-022-01085-0>. 2, 3, 10, 11, 12, 14

A Supplementary Appendix

This project used computational tools which are fully open-source. Any comments or suggestions may be sent to me at seh325@cornell.edu, or raised as an issue on the Github project.

Data analysis with the *Tidyverse* (Wickham et al. 2019) package, for the R language (R Core Team 2023), made the empirical analysis for this paper possible.

A.1 UK Biobank Dataset Construction

A.1.1 Identifying Family Clusters by Genetic Similarity

The UKB does not directly identify participants’ family members, so I define family links using the genetic relatedness data. I follow the approach in Muslimova et al. (2025), classifying genetically related pairs using the kinship coefficient and the identity-by-state zero coefficient (IBS₀). These thresholds come from the KING relationship-inference procedure used to construct the UKB kinship matrix, which identifies genetically related pairs up to third-degree relatives (Manichaikul et al. 2010). The kinship coefficient measures the degree of genetic relatedness between two individuals, while IBS₀ measures the fraction of genotyped markers for which the two individuals share no alleles. These two measures jointly distinguish different forms of close genetic relatedness.

Table A1: Genetically Identified Family Links in the UKB.

Genetic link	Kinship Coefficient	IBS ₀	Count
Duplicate / Monozygotic twins	> 0.3540		179
1st degree / Parent-child	0.1770–0.3540	< 0.0012	6,249
1st degree siblings	0.1770–0.3540	≥ 0.0012	22,614
2nd–3rd degree relatives / cousins	0.0442–0.1770		77,878
Total			106,920

I classify duplicate observations and monozygotic twins as pairs with kinship coefficient above 0.3540. First-degree relatives are pairs with kinship coefficient between 0.1770 and 0.3540. Within this first-degree interval, parent-child pairs and full siblings are distinguished by IBS₀. Parent-child pairs have IBS₀ below 0.0012, since a biological parent and child should share one allele at nearly every locus. Full siblings have IBS₀ weakly above 0.0012, since full siblings may share zero, one, or two alleles at a given locus. Second- and third-degree relatives are defined as pairs with kinship coefficient between 0.0442 and 0.1770.

Table A1 reports the number of genetically related pairs in each category, among all pairs of individuals with kinship coefficient at least 0.0442, corresponding to cousins or closer

relatives. The analysis sample includes individuals with at least one genetically identified first-degree connection in the UKB relatedness file, and filtered to those with non-missing data in occupation codes and education variables. The counts differ from [Muslimova et al. \(2025\)](#) thanks to this filtering.

I use these pairwise genetic links to construct family identifiers. Each UKB participant is treated as a node in a family-relatedness network, and each genetically identified relationship is treated as an edge between two participants. Family identifiers are then assigned from the connected components of this network. The primary analysis sample is based on individuals with at least one genetically identified full sibling in the UKB, since full-sibling links are the relevant family relationship for the within-family design and the sibling-based parental PGI imputation procedure.

The design of the UKB allows for a larger number of family clusters than would be expected from a random sample of the British population. This is because study enrolment letters were sent to households, and the household head (or to whomever the letter was addressed) could optionally enrol their family members in addition to themselves.

A.1.2 Occupation Wage Imputation


The UKB does not report continuous individual wages, but it does record labour-market information that can be used to impute wages, including 4-digit Standard Occupational Classification (SOC) codes and reported working hours. I construct occupation-based wages using the supplementary files and replication code from [Kweon et al. \(2025\)](#), following the same imputation procedure as [Carvalho \(2025\)](#). The resulting measure should be interpreted as an occupation-based prediction of labour-market earnings, rather than a directly observed individual wage.

The imputation procedure combines information from three UK data sources. First, the Labour Force Survey (LFS) provides individual-level wage data and labour-market covariates. Second, the Annual Survey of Hours and Earnings (ASHE) provides high-quality occupation-level wage information by SOC code. Third, the British Household Panel Survey (BHPS) is used to validate the imputation procedure in an independent dataset. The procedure estimates a parametric prediction model for log wages in the LFS, using ASHE mean and median wages for each occupation group together with demographic and labour-market predictors. The estimated model is then applied to UKB participants using their SOC code and observed labour-market characteristics.

The key predictor is the participant’s 4-digit SOC code. Rather than including a full set of 4-digit occupation fixed effects, which would create sparse cells after interacting occupation with age, sex, full-time status, and year, the [Kweon et al. \(2025\)](#) procedure uses ASHE mean and median wages for each occupation group as continuous predictors. This preserves detailed occupational wage information while avoiding failed or noisy imputations in small occupation-demographic cells. I use the imputed hourly wage measure, and construct annual occupation income by multiplying the imputed hourly wage by reported weekly hours and 52 weeks.

A.2 Genetic Effects

Figure A1: Correlation for Ed PGI, and Ed PGI Minus Parental Mean, with Demographic Variables.



Note: Multivariate regression estimates for the linear model $Z_i = \phi' \mathbf{X}_i + \eta_i$, where $\mathbf{X}_i$ is the vector demographic variables described by the labels on the y -axis. The dots are the point estimates, and bars are the 95% confidence intervals. The random component refers to the Ed PGI after residualizing out parental values,

A.3 Causal Mediation (CM) Decomposition

The decomposition of the average partial effect into direct and indirect components follows from a straightforward application of the multivariate chain rule. Define:

$$g(a, b) = \mathbb{E} [Y_i(a, D_i(b))].$$

The average partial effect, direct effect, and indirect effect can then be written as:

$$\begin{aligned}\text{Average partial effect} &= \frac{d}{dz}g(z, z), \\ \text{Direct effect} &= \frac{\partial}{\partial a}g(a, Z_i)\Big|_{a=z}, \\ \text{Indirect effect} &= \frac{\partial}{\partial b}g(Z_i, b)\Big|_{b=z}.\end{aligned}$$

The average partial effect moves both arguments of g simultaneously with z , whereas the direct and indirect effects move only the first and second arguments respectively, holding the other fixed. By the multivariate chain rule:


$$\frac{d}{dz}g(z, z) = \frac{\partial g}{\partial a}(z, z) + \frac{\partial g}{\partial b}(z, z),$$

which establishes that the average partial effect is exactly the sum of the direct and indirect effects, evaluated at $z = Z_i^{\text{Child}}$. This result is exact under the smoothness and regularity conditions assumed in [Section 3](#), with no further approximation required.

A.4 Returns to Education by the MLSA Rise

The UK Biobank includes a total of 500,000 members of the British public, with a roughly even distribution of birth years over 1945–1965. In 1972, the British government increased the Minimum School Leaving Age (MSLA) by one year, from age 15 to 16. The change took effect at the start of the 1972-73 school year, and applied to those born on or after 1 September 1957. The reform therefore induced extra school years for individuals born around date 1 September 1957. British school children sit exams for the GCSE qualification at age 16, so the reform pushed a larger share of the UK Biobank cohorts to remain in school to 16 and to complete these qualifications.

The 1972 MLSA change has been previously used for labour market education returns ([Oreopoulos 2006](#)), and (null) mortality effects of compulsory education ([Clark & Royer 2013](#)). The UK Biobank population overlaps well with the birth date cutoff, and so has been used in a series of papers for estimating education returns in the UK Biobank, for health outcomes ([Barcellos et al. 2023, 2025, Davies et al. 2018](#)) and occupation imputed income ([Barcellos et al. 2021](#)).

Figure A2: Mean Effects of the MSLA Rise.

Note: These figures show the effects of the 1972 MSLA rise on multiple outcomes, where individuals born after September 1957 (the dashed line) were legally compelled to stay in school until age 16, whereas before only until age 15. Each dot is the mean for UK Biobank respondents with non-missing education data born in that month-year, and the lines are local polynomials either side of September 1957.

Birth month is the running variable in a fuzzy Regression Discontinuity Design (RDD), using the indicator for being born September 1957 and after as an instrument. Because not all those eligible necessarily stayed beyond 15, the post-cutoff indicator instruments education variables, where mentioned. I estimate the following RDD estimands, estimating the effect of the MSLA change on education variables — represented by D_i .

$$\lim_{v \rightarrow -1957.75} \mathbb{E}[D_i | V_i = v], \quad \lim_{v \rightarrow +1957.75} \mathbb{E}[D_i | V_i = v]$$

where V_i represents individual i 's birth year plus birth month divided by 12 (so that 1957.75 represents born in September 1957), and $\rightarrow^-$ refers to the limit from below, $\rightarrow^+$ above. The goal is to identify is the average compliance rate of the MSLA rise,

$$\mathbb{E}[D_i(1) - D_i(0)] = \lim_{v \rightarrow +1957.75} \mathbb{E}[D_i | V_i = v] - \lim_{v \rightarrow -1957.75} \mathbb{E}[D_i | V_i = v].$$

The MSLA rise significantly increases education years, completing of age 16 secondary school qualifications (GCSEs) — these are measured of education coded by qualifications, as used by [Okbay et al. \(2022\)](#) in calculating the Ed PGI. It had an even larger effect on age at leaving school, which is the basis of the education variables that [Barcellos et al. \(2021\)](#) focuses on. The mean of these variables by birth month are shown in [Figure A2](#). Note that the economics literature has argued that null estimates for education returns in the MSLA design are a result of the MSLA rise inducing a low-quality margin of education ([Clark & Royer 2013](#)), so while age leaving education variables are more significant for this instrument they may not be economically relevant.

I also use the RDD design to measure changes in later labour market outcomes. The UK Biobank asks questions of the sector people work in, giving high detail occupation codes (i.e., SOC codes) Using this value, and other demographic details, I impute occupation coded income for the UK Biobank respondents using data from the British Annual Survey of Hours and Earnings.¹⁷ This gives occupation coded hourly income, and annual wage calculated from hours worked as outcome measures in the MSLA rise, affected by the MSLA rise

$$\lim_{v \rightarrow +1957.75} \mathbb{E}[Y_i | V_i = v] - \lim_{v \rightarrow -1957.75} \mathbb{E}[Y_i | V_i = v].$$

The final panels of [Figure A2](#) show these outcomes, and this reduced-form effect of the MSLA is not significantly different from zero. [Barcellos et al. \(2021\)](#) analyse this same RDD with UKB data.

¹⁷Thanks to [Kweon et al. \(2025\)](#) for sharing code on this imputation.


Table A2: Estimates of the MLSA Change Effects.

Panel A: First-stage.	First-stage.				Reduced form.			
	GCSE completion		Education years		Hourly wage		Annual income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Outcome:							
MSLA Effect	0.089 (0.006)	0.247 (0.011)	0.158 (0.060)	0.314 (0.058)	0.006 (0.007)	0.006 (0.007)	0.010 (0.010)	0.010 (0.010)
<i>F</i> -statistics	216	530	6.85	29.7	0.939	0.939	0.830	0.830
CCT bandwidth	2.66	1.60	3.21	2.60	4.10	4.10	3.87	3.87
Bandwidth observations	64,011	24,185	77,881	38,718	99,646	99,646	93,885	93,885
Panel B: Second-stage.	OLS.				MLSA IV.			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome: Log Occupation Hourly Wage								
GCSE completion	0.372 (0.008)	0.193 (0.008)			0.070 (0.094)	0.068 (0.039)		
Education years			0.061 (0.000)	0.034 (0.002)			0.044 (0.047)	0.057 (0.029)
CCT bandwidth	10.0	10.0	10.0	10.0	2.54	2.71	4.21	4.08
Bandwidth observations	245,740	149,732	245,740	149,732	62,027	39,938	101,570	58,987
Outcome: Log Occupation Annual Income								
GCSE completion	0.392 (0.015)	0.188 (0.009)			0.100 (0.151)	0.046 (0.065)		
Education years			0.065 (0.001)	0.032 (0.002)			0.062 (0.080)	0.048 (0.049)
CCT bandwidth	10.0	10.0	10.0	10.0	2.61	2.86	3.84	4.16
Bandwidth observations	245,740	149,732	245,740	149,732	64,011	42,307	93,885	60,155
Qualification definition	Yes		Yes		Yes		Yes	
School age definition		Yes		Yes		Yes		Yes

Note: These point estimates are calculated among all UK Biobank participants born within a bandwidth of the MLSA rise cutoff (not just the sibling subsample). The odd columns use a definition of education attainment (age 16 schooling) based on reported education qualifications; even columns report based on school leaving age. The Panel B OLS results show estimates of an OLS specification for the correlation between education and wages/income with local polynomial trends each side of 1957.75 birthdate, attempting to mimic an RD set-up with a hard-coded bandwidth of 10 years for comparison.

I estimate this RDD equation with triangular kernel weights (weighting closer to the cutoff), an optimal bandwidth (Calonico et al. 2014, 2015), and non-parametric controls on either side of the cutoff (Gelman & Imbens 2019). These estimates are collected in Table A2.

Figure A3: Fuzzy RDD Estimates, Sensitivity to Bandwidth Choice.



Note: These figures show the sensitivity of fuzzy RD estimates using the MLSA discontinuity as an IV for education years (qualification measure) or age leaving school, using a different bandwidth choice for each. The shaded regions are the 95% confidence interval. The dotted line with CTT refers to the Calonico et al. (2014) optimal bandwidth choice, calculated using *rdrobust* in *R* (Calonico et al. 2015).

The MLSA rise increased years of education, as measured in both qualification definitions and school leaving age definition. However, the outcome of occupation coded wages and annual income are not strong; correspondingly, the IV estimates using the MLSA rise for education returns are not distinguishable from zero. This holds true for a range of bandwidth choices, shown in Figure A3.

A.5 Instrumental Variables Connection

The CM analysis in this paper is conceptually similar to an instrumental variables analysis; the main difference is that it does not assume the exclusion restriction, and instead attempts to model and estimate this direct effect. One might consider using the Ed PGI as an instrument for education years, interpreting the reduced-form effect of the Ed PGI on earnings only through its first-stage effect on education. This is the logic of Mendelian Randomisation (MR), where genetic variants are used as instruments for health conditions, and any direct effect of the genetic index on the outcome is treated as a violation of the exclusion restriction.

Widding-Havneraas et al. (2026), for example, uses the Ed PGI as an instrument for education in an MR design; this approach is not well suited to the present question as it assumes the direct effects are exactly zero for every individual, giving estimates for returns

to education contaminated by plausibly non-zero direct effects. The CM approach instead estimates both the direct genetic effect and the indirect education effect.

There is a literature for modelling such exclusion violations in economics ([Conley et al. 2012](#), [Kolesár et al. 2015](#)), and in genetic settings using alternative assumptions tailored to MR ([Bowden et al. 2015](#), [Spiller et al. 2019](#)). These alternative approaches require structure on how invalid instruments violate exclusion; the most common adjustment for genetic settings, MR-egger ([Bowden et al. 2015](#)), relies on the direct genetic effects being uncorrelated with instrument strength, paralleling the invalid instruments assumption in the econometrics literature ([Kolesár et al. 2015](#)). This is difficult to justify in the present setting, where the SNPs that most strongly predict education could have correspondingly larger with direct effects on earnings through cognitive skills, non-cognitive traits, health, field of study, occupational sorting, or other labour-market relevant channels.

A.6 Returns to Education Literature

Table A3: Returns to Education, British Estimates from the Literature.

OLS	IV	Year	Reference
29.2	5.7	2010	Dickson, M. and Smith, S., 2011. What determines the return to education: an extra year or a hurdle cleared?. <i>Economics of education review</i> , 30(6), pp.1167-1176.
12.2	6.6	2009	Devereux, P.J. and Fan, W., 2011. Earnings returns to the British education expansion. <i>Economics of Education Review</i> , 30(6), pp.1153-1166.
9.6	5.3	2009	Devereux, P.J. and Fan, W., 2011. Earnings returns to the British education expansion. <i>Economics of Education Review</i> , 30(6), pp.1153-1166.
9.5	6.9	2006	Grenet, J. 2013. Is it enough to increase compulsory education to raise earnings? Evidence from French and British compulsory schooling laws. <i>Scandinavian Journal of Economics</i> 115: 176-210.
7.8	6.2	2009	Devereux, P.J. and Fan, W., 2011. Earnings returns to the British education expansion. <i>Economics of Education Review</i> , 30(6), pp.1153-1166.
6.0	4.4	1999	Blanchflower, D. and Elias, P., 1999. Ability, schooling and earnings: Are twins different?. Unpublished manuscript, Warwick University.
6.9	6.6	2009	Devereux, P.J. and Fan, W., 2011. Earnings returns to the British education expansion. <i>Economics of Education Review</i> , 30(6), pp.1153-1166.
7.0	7.0	2001	Devereux, P.J. and R.A. Hart. 2010. Forced to be Rich? Returns to Compulsory Schooling in Britain. <i>Economic Journal</i> 120(549): 1345-64.
7.3	7.7	1999	Bonjour, D., Cherkas, L.F., Haskel, J.E., Hawkes, D.D. and Spector, T.D., 2003. Returns to education: Evidence from UK twins. <i>American Economic Review</i> , 93(5), pp.1799-1812.
4.8	5.5	1991	Dearden, Lorraine. 1999. Qualifications and earnings in Britain: how reliable are conventional OLS estimates of the returns to education? IFS Working Papers W99/07, Institute for Fiscal Studies.
4.8	5.5	1981	Dearden, L. (1998), "Ability, Families, Education and Earnings in Britain," Institute for Fiscal Studies Working Paper No. W98/14.
0.0	1.2	1947	Chib, Siddhartha, and Liana Jacobi. 2016. "Bayesian Fuzzy Regression Discontinuity Analysis and Returns to Compulsory Schooling." <i>Journal of Applied Econometrics</i> 31: 1026-1047.
5.0	9.9	1970s	Harmon, C. and Walker, I. (2000), "Returns to the Quantity and Quality of Education: Evidence for Men in England and Wales," <i>Economica</i> , 67: 19-35.
4.6	10.2	1972	Dickson, Matt. 2013. "The Causal Effect of Education on Wages Revisited." <i>Oxford Bulletin of Economics and Statistics</i> 75 (4): 477-498.
7.8	14.7	1947	Oreopoulos, P. 2006. Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter <i>American Economic Review</i> 96(1) 152-175
4.9	14.0	1992	Harmon, C. and Walker, I. (1999), "The Marginal and Average Return to Schooling in the UK," <i>European Economic Review</i> , 43(4-6), pp.879-887.
6.1	15.3	1947	Harmon, C. and I. Walker. 1995. Estimates of the economic return to schooling for the United Kingdom. <i>American Economic Review</i> 85: 1278-86.

Note: This table shows the education returns estimates on later-life labour market earnings in the economics literature, using an IV to causally estimate the effect of a year of education. Literature review kindly provided by [Patrinos & Psacharopoulos \(2026\)](#).