

NBER WORKING PAPER SERIES

THE EFFECT OF EDUCATION ON THE RELATIONSHIP BETWEEN GENETICS,
EARLY-LIFE DISADVANTAGES, AND LATER-LIFE SES

Silvia H. Barcellos
Leandro Carvalho
Patrick Turley

Working Paper 28750
<http://www.nber.org/papers/w28750>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2021, Revised September 2021

Jiewen Luo and Andrew Yu provided excellent research assistance. This paper benefited from discussions with Dan Benjamin, Pietro Biroli, Dalton Conley, David Cesarini, Ben Domingue, Jason Fletcher, Rachel Griffith, Mireille Jacobson, Phil Koellinger, Adriana Lleras-Muney, Nick Papageorge, Tuomas Pekkarinen, Lauren Schmitz, and from the feedback of seminar participants at Boston University, Lehigh University, Stanford University, University of Rome Tor Vergata, UCLA, USC and at the 2019 All-CA Labor Economics Conference at UCSC. We thank Hyeokmoon Kweon and Phil Koellinger for help with the ASHE data and Alex Young for providing the imputed polygenic indices of parents. Research reported in this publication was supported through the Russell Sage Foundation grant 98-16-16, NIA grants K99-AG062787-01, K01AG050811, 1K01AG066999-01A1, R21AG060447, RF1AG055654, 3RF1AG055654-01A1S1, and R56AG058726, Open Philanthropy grant 010623-00001, and by the USC Population Research Center. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. This research has been conducted using the UK Biobank Resource under Application Number 15666. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Silvia H. Barcellos, Leandro Carvalho, and Patrick Turley. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Effect of Education on the Relationship between Genetics, Early-Life Disadvantages,
and Later-Life SES

Silvia H. Barcellos, Leandro Carvalho, and Patrick Turley

NBER Working Paper No. 28750

May 2021, Revised September 2021

JEL No. I24,I26,J31

ABSTRACT

This paper investigates whether education weakens the relationship between early-life disadvantages and later-life SES. We use three proxies for advantage that we show are independently associated with SES in middle-age. Besides early, favorable family and neighborhood conditions, we argue that the genes a child inherits also represent a source of advantages. Using a regression discontinuity design and data for over 110,000 individuals, we study a compulsory schooling reform in the UK that generated exogenous variation in schooling. While the reform succeeded in reducing educational disparities, it did not weaken the relationship between early-life disadvantages and wages. This implies that advantaged children had higher returns to schooling. We exploit family-based random genetic variation and find no evidence that these higher returns were driven by genetically-influenced individual characteristics such as innate ability or skills.

Silvia H. Barcellos
Center for Economic and Social Research
University of Southern California
635 Downey Way
Los Angeles, CA 90089
and NBER
sbarcell@usc.edu

Patrick Turley
Center for Economic and Social Research
University of Southern California
635 Downey Way
Los Angeles, CA 90089
pturley@usc.edu

Leandro Carvalho
Center for Economic and Social Research
University of Southern California
635 Downey Way
Los Angeles, CA 90089
leandro.carvalho@usc.edu

A data appendix is available at <http://www.nber.org/data-appendix/w28750>

“Education then, beyond all other devices of human origin, is a great equalizer of the conditions of men – the balance wheel of the social machinery.”

Horace Mann, pioneering American educator, 1848

“In America, education is still the great equalizer.”

Arne Duncan, U.S. Secretary of Education, 2011

1. Introduction

Children born to richer parents are more likely to achieve economic success in adulthood (Jencks 1979; Solon 1999; Almond and Currie 2011). Education is considered by many to be the “great equalizer” and education policy a tool to “level the playing field,” ensuring that kids from different backgrounds have similar opportunities for success. Others contend that, while education may reduce poverty, it might be less effective in promoting mobility because children born to richer parents or with higher ability have higher returns to schooling (Cameron and Heckman 2001) and consequently will choose to invest more in education.

This paper investigates whether education weakens the relationship between early-life advantages and later-life socioeconomic status (SES). Using UK Biobank data for over 110,000 individuals, we examine how these relationships were affected by a compulsory schooling change that generated exogenous variation in schooling.

We consider early-life environmental advantages at the *family* and at the *neighborhood* level: being born in a smaller family² and in a neighborhood with higher average education³. We also consider the genes that one inherits as *individual-level* advantages. Using genetic data, we construct a polygenic index (PGI) for educational attainment (EA), which has been shown to predict a host of SES outcomes (Lee et al. 2018; Belsky et al. 2018). Because of Mendelian inheritance, one’s PGI is randomly assigned conditional on the PGI of the parents. We exploit this variation to study a gene-by-environment interaction: how a compulsory schooling change can modify how much of an advantage is conferred by one’s luck in the “genetic lottery” (Kweon et al. 2020).

² Following the large empirical literature on the trade-off between quantity and quality of children (e.g. Steelman et al. 2002, Chan et al. 2019).

³ More specifically, the fraction of the adult population in the respondent’s birth town around her birth date who stayed in school until at least age 15. Recent work shows that the quality of the neighborhood in which a child grows up affects her later-life SES (e.g. Chetty et al. 2016; Chetty and Hendren 2018a, 2018b).

Consistent with the literature on the importance of initial endowments (e.g., Lee et al. 2018; Black et al. 2005; Mogstad and Wiswall 2016; Chetty and Hendren 2018a, 2018b), our three proxies for early-life advantages are good predictors of education and economic success in adulthood. While each proxy is measured at a distinct level, it can capture in principle variation from the other levels (e.g., since genes are inherited from one’s parents, the PGI captures both individual-level and family-level advantages). However, the predictive power of each proxy is remarkably similar if we estimate a joint model where the three proxies are entered together, indicating that each proxy captures an independent dimension of disadvantage.

While these associations do not represent causal relationships (e.g., Kong et al. 2018; Angrist et al. 2010), each proxy is fixed early in life, suggesting that they potentially capture some other early-life, causal factors (Rietveld et al. 2013; Okbay et al. 2016; Harden et al. 2020). Indeed, we exploit the random genetic variation mentioned above to show that the EA PGI itself has causal effects on SES outcomes.⁴

The compulsory schooling change we study, the Raising of the School Leaving Age (ROSLA) Order of 1972, increased the minimum age at which students in England, Scotland, and Wales could drop out of school from 15 to 16 years (Oreopoulos 2006; Clark and Royer 2013; Barcellos, Carvalho and Turley 2019). The reform only affected students born after September 1, 1957, generating a discontinuity in the relationship between education and date of birth that can be exploited using a regression discontinuity design. We use data from the UK Biobank ($N = 110,000$) to study the effects on schooling, income, and occupational wages 34-38 years after the reform. We estimate that an additional year of schooling increased occupational wages by 5% on average—comparable to previous estimates using the 1972 ROSLA (Grenet 2013).

The ROSLA was successful in reducing educational disparities. For example, years of schooling increased by 0.39 years for those in the most disadvantaged tercile (hereafter, the bottom tercile) of the distribution of the EA PGI and by 0.27 years for those in the most advantaged tercile (hereafter, the top tercile) – this is also true for the other two proxies for advantages. We find similar effects on education qualifications (namely the CSE and O-Level examinations, which are taken at age 16) that are valued in the labor market. These results suggest that the ROSLA may have effectively narrowed the gap in skills between children from different backgrounds.

⁴ This has been shown in previous studies (e.g., Belsky et al. 2018; Lee et al. 2018; Kweon et al. 2020).

However, these educational gains did not translate into reductions in economic disparities. While the ROSLA increased wages across the board, it did not close the wage gap. If anything, the reform may have widened it. For example, the reduced-form estimates imply that the ROSLA increased the wages of children in the top tercile of the EA PGI distribution by 1.6 percentage points more than it increased the wages of those in the bottom tercile, with similar patterns for family size and neighborhood SES. Using two-stage least squares, we estimate that the return to an additional year of schooling for those in the top terciles of the EA PGI, family size, and neighborhood SES distributions was 6, 4, and 4 percentage points higher, respectively, than the returns of those in the bottom terciles of these distributions. Using a marginal treatment effects framework (Carneiro et al. 2011; Brinch et al. 2017; Kowalski *forthcoming*), we show that these results cannot be explained by differences in unobservable characteristics of ROSLA between advantaged and disadvantaged compliers.

These results raise the question of why children from advantageous backgrounds benefited more from the additional schooling. Some economists would argue differences in returns to schooling may be justified if they reflect differences in innate ability, skills, or talent (Loury 1981). Many may also contend that it is unfair if children from advantageous backgrounds have higher returns due to, for example, their parents' connections (Bénabou 2000).

We exploit family-based random genetic inheritance to investigate whether the gradient in the returns to schooling is driven by genetically-influenced individual characteristics. By conditioning on the EA PGI of parents, we can estimate the causal effect of the child's EA PGI.⁵ The results suggest that a *ceteris paribus* increase in the child's EA PGI causes a reduction in the returns to schooling. In contrast, having parents with higher EA PGIs is associated with higher returns to schooling. These results do not support the hypothesis that genetically-influenced individual characteristics, which could include innate ability or skills, are responsible for the positive gradient in returns to schooling. They *suggest* instead that children from higher-SES backgrounds had higher returns because of environmental advantages (e.g., having more educated parents, higher family income, or growing up in a richer neighborhood) that may confer unequal opportunities. We note, however, that this particular analysis is not as well-powered as the other analyses in the

⁵ In Figure 3 below, we show that once we condition on the EA PGI of parents, the EA PGI of children is orthogonal to a number of predetermined variables. This "balance test" suggests that our method is successful in isolating random genetic inheritance.

paper because identification only comes from the 10% of the sample for whom we have data from a parent or sibling (see section 4.D).

Our paper makes a number of contributions. First, we bring together two distinct literatures: a literature in economics on childhood circumstances and adulthood SES (e.g. Almond et al. 2018) and a literature in social genomics on genotypes and later-life outcomes (Barth et al. 2020; Lee et al. 2018; Belsky et al. 2018). We document that inheriting certain genes, being born into a smaller family, and in a higher-SES neighborhood capture independent dimensions of advantage. For economists, the EA PGI has a number of appealing characteristics as a proxy for advantage. Not only it is objectively measured, it is also determined at conception and invariant to environmental conditions. While the EA PGI may be correlated with one's environment, we illustrate how it is possible to decompose it into individual-level genetic variation and variation in environmental conditions using the EA PGIs of parents (Domingue and Fletcher 2020; Young et al. 2020). More importantly, this individual-level genetic variation is random, providing a source of exogenous variation in early-life advantage.

Second, we study how differences in the *quantity of schooling* and in the *returns to schooling* separately contribute to intergenerational mobility. Previous work on the effects of education on intergenerational mobility was unable to separate the two because the educational reforms studied changed not only the quantity of education but also key elements of the educational system (Meghir and Palme 2005; Pekkarinen et al. 2009; Aakvik et al. 2010; Andreoli et al. 2020; Bertrand, Mogstad and Mountjoy 2020). Our results underlie the importance of making this distinction: in our context advantaged children have substantially higher returns to schooling, which may present a challenge to increasing mobility (Solon 2004).

Finally, we are among the first papers to combine quasi-random variation in environment with random genetic variation to study how genetics interacts with a key environmental factor – education – to shape economic success. Our findings speak to three distinct literatures in economics: (a) on the importance of genetic inheritance for the intergenerational transmission of SES (Bowles and Gintis 2002; Black et al. 2020; Sacerdote 2011; Fagereng et al. 2021); (b) on gene-by-environment interactions (Papageorge and Thom 2020; Biroli and Zünd 2020)⁶; and (c) on heterogenous returns to schooling (Card 2001).

⁶ There is a large literature in social genetics on gene-by-environment interactions (e.g., Caspi et al. 2002; Caspi et al. 2003; Shanahan et al. 2008). Most studies in this literature are correlational.

We advance the first literature (a) by showing that educational policy has the potential to affect the contribution of genetic inheritance for the intergenerational transmission of SES. Our work builds on studies in the second literature (b) which have furthered our understanding of how genes and environment interact by exploiting natural experiments (Schmitz and Conley 2016, 2017; Barcellos et al. 2018; Biroli and Zwysig 2021) or by using within-family designs (Domingue and Fletcher 2020; Morris et al. 2020; Brumpton et al. 2020; D’onofrio et al. 2013). We take a step forward by combining a natural experiment *with* a family design (Fletcher and Conley 2013). Our contribution to the third literature (c) is to study how genes influence the returns to schooling. The literature has tried to distinguish between heterogeneous returns by family SES (Altonji and Dunn 1996; Deschênes 2007) and heterogeneous returns by ability (Ashenfelter and Rouse 1998; Heckman et al. 2006, Heckman et al. 2018). One challenge is that commonly-used measures of ability (e.g., AFQT) may reflect family SES. By using random genetic variation determined at conception, we are able to disentangle genetically-influenced differences in returns to schooling from differences in returns due to environmental advantages. To the extent that EA PGI captures ability, we find no evidence that the advantaged children in our setting had higher returns to schooling because they had greater innate ability.

The rest of the paper is structured as follows. In Section 2, we give more details about the data and introduce our proxies for early disadvantages. Section 3 discusses the interpretation of the genetic proxy (i.e., the EA PGI) and the method used to isolate the random genetic inheritance. Section 4 examines how the compulsory schooling law affected the relationship between early-life disadvantages and later-life SES. Section 5 investigates whether differences in unobservable characteristics between advantaged and disadvantaged compliers can explain the results presented in Section 4. Robustness exercises are presented in Section 6. Section 7 concludes.

2. Data and Proxies for Early-Life Advantages

We use data from the UK Biobank, a large, population-based prospective study initiated by the UK National Health Service (NHS) (Sudlow et al. 2015). More than half a million individuals ages 40 to 69 were assessed between 2006 and 2010 in 22 assessment centers distributed throughout

the UK (Allen et al. 2012) – see Appendix K.⁷ The assessment included a self-completed touchscreen questionnaire, an in-person interview, physical measurements, and the collection of biological samples. The self-completed questionnaire collected data on, among other things, country of birth and year of immigration; qualifications and school-leaving age; household income; and family history. Information that was not collected via the self-completed touchscreen questionnaire, such as occupation and residence at birth, was collected in a subsequent computer-assisted personal interview (CAPI). Study participants were also genotyped using blood samples collected at the end of the assessment visit.

Our main outcomes of interest will be the participant’s age at the time she left school; whether she passed qualification exams that are taken at age 16 – namely the Certificate of Secondary Education (CSE) and the General Certificate of Education (GCE) Ordinary Level (also known as the O-level); annual household income (reported in five brackets); and wages imputed based on one’s occupation. During the interview, respondents answered a series of questions about their job (or last job if respondent had retired recently), which the interviewer used to classify the respondent’s occupation among more than 400 detailed categories. We use the 2009 Annual Survey of Hours and Earnings (ASHE) to match such categories to median wages for each occupation (Kweon et al. 2020).⁸ We view this measure of occupational wages as complementary to our measure of household income: on one hand, income was reported in brackets and occupational wages is a better measure of permanent individual income; on the other, household income is available for a larger fraction of the sample.⁹ We find that both measures yield qualitatively similar results. Appendix C shows that the two measures are closely related.

A. Proxies for Early-Life Advantages

We use measures to proxy for early-life advantages at three levels: neighborhood, family, and individual. A recent literature shows that children who grow up in higher SES neighborhoods have better socioeconomic outcomes, including earnings and college attendance (Chetty and

⁷ Although the UK Biobank is not nationally representative, our estimates have internal validity because there is no differential selection on the two sides of the September 1, 1957 cutoff – see Appendix A.

⁸ We used gender-specific wages that were calculated among full-time employees and included all labor earnings (such as bonus, tips, overtime, etc.).

⁹ Occupational wage is missing for 26.8% of UKB respondents while household income is missing for 14.5%. We show in the Robustness section that results are qualitatively similar if we use income instead of wages. See also Appendix H.

Hendren 2018a, 2018b). As a proxy for neighborhood SES, we use a measure of average education in the individual’s birthplace. UKB participants reported the town or district where they first lived when they were born. The coordinates of this locality were used to identify the local district (as of 1961) where the respondent resided at that time. The measure of neighborhood SES corresponds to the fraction of adults in the local district who stayed in school until at least age 15 according to the 1961 Census.^{10,11,12} There are 1,436 different neighborhoods in our sample with an average population of about 35,000 and a median population of about 18,000.¹³

We use family size as a proxy for family-level advantages, following the large empirical literature on the trade-off between quantity and quality of children (e.g. Steelman et al. 2002; Chan et al. 2019). Our measure of family size corresponds to the numbers of full brothers and full sisters (i.e., excluding half-siblings, step-siblings, or adopted siblings) reported by the study participant, including those who had already died. Because individuals with fewer siblings tend to have higher SES later in life, we reverse code family size (smaller families have a higher proxy value) so it is positively correlated with SES – as it is the case with the other proxies.

As our measure of individual-level advantages, we use a polygenic index¹⁴ (PGI) for educational attainment (EA). A PGI is an individual-level predictor constructed from up to millions of genetic markers. Below, we describe how we constructed the EA PGI and discuss its interpretation.

Human DNA is made up of twenty-three pairs of long molecules, called *chromosomes*. While any pair of individuals are identical for 99.8% of their DNA, there are tens of millions of locations in the genome where individuals differ by a single genetic marker. These locations are called *Single Nucleotide Polymorphisms* (SNPs). At the vast majority of SNPs, people can have one of two possible genetic variants. The variant that a person has is called their *allele*. In genetic data, one of the two possible alleles is arbitrarily chosen as the *reference allele*. Because individuals have two copies of each chromosome, they will either have 0, 1, or 2 copies of the

¹⁰ We chose the threshold of age 15 because it maximized the variation in neighborhood characteristics. Importantly, this variable was chosen before assessing its association with our measures of SES and education. It captures a large fraction of the potential neighborhood variation at this level of aggregation.

¹¹ These data were generously made available by Vision of Britain (www.visionofbritain.org.uk).

¹² For Scotland, these figures were calculated at the county level using data from the 1951 Census.

¹³ Chetty and Hendren (2018a) use commuting zones (CZ) as a measure of neighborhood. There are 741 different CZs in the United State with an average population of about 380,000 per CZ.

¹⁴ Polygenic indexes are also often called “polygenic scores,” “polygenic risk scores,” or “genetic risk scores” in the literature. All of these terms correspond to the same thing.

reference allele. The number of reference alleles that an individual has at a SNP is called their *genotype* for that SNP.

A PGI is constructed using estimates from a *Genome Wide Association Study* (GWAS). A GWAS scans the entire genome and estimates associations between individual genotypes and an outcome of interest. Specifically, a GWAS is a series of regressions of some outcome onto the genotype of each SNP, one at a time, and a set of covariates which normally include sex, age, and the first several principal components of the genetic data.¹⁵

A *polygenic index* (PGI) is a weighted sum of SNP genotypes:

$$PGI_i = \sum_j g_{ij} w_j, \quad (1)$$

where PGI_i is the polygenic index for individual i ; $g_{ij} \in \{0,1,2\}$ is individual i 's genotype at SNP j ; and w_j is the weight for SNP j .¹⁶ The weights in a PGI are derived from coefficients estimated in a GWAS.¹⁷ In this paper, we use GWAS coefficients based on Lee et al. (2018), the largest currently available GWAS for educational attainment (EA).¹⁸ The PGI based on this GWAS can explain 12-13% of the variation in EA (Lee et al., 2018).¹⁹

We note that each of our three measures of advantage may proxy for circumstances at other levels. Also, with the exception of family size, the proxies are fixed at birth, which means they could not have been affected by the compulsory schooling law that we will study. We show in Appendix A that the reform did not affect family size.

We keep in our sample individuals for whom at least one proxy is available.²⁰ About ninety percent of them have data on all three proxies. In the regression results below, we include an

¹⁵ These principle components are included to account for ancestry-related omitted variable bias (see Section 3.A).

¹⁶ GWASs for educational attainment have shown that education is associated to a large number of genetic markers, each with a small influence. PGIs are a way to aggregate these many influences and construct a genetic marker that is sufficiently predictive to use in empirical applications.

¹⁷ There are several methods for producing PGI weights from GWAS coefficients, but each of them transforms the GWAS coefficients in a way that is meant to account for the correlation structure that exists in the genome. We use a Bayesian method called LDpred (Vilhjálmsson et al. 2015).

¹⁸ We use a version of the GWAS coefficients from Lee et al. (2018) that omit the subset of individuals from the UK Biobank that we use in our analyses. Specifically, we conduct a GWAS in the UK Biobank that exactly matches the specification in Lee et al., except that we exclude individuals who were born within 10 years of September 1, 1957 and their family members. This is done to avoid overfitting.

¹⁹ Because Lee et al. (2018) is a GWAS of level of EA, our EA PGI is calibrated to predict levels of EA. Johnson et al. (2020) argue that for interaction studies such as ours, it may be more appropriate to use a GWAS of the variability of the outcome rather than the level. Nevertheless, we use PGI based on levels because Lee et al. has a much larger sample than any available GWAS on the variability of EA. Using a smaller sample would greatly reduce the power of our study.

²⁰ Only 77 UKB participants lacked all three proxies.

indicator (and interaction terms when relevant) for whether the individual is missing each proxy. Our results are very similar if we restrict to only those with complete data. The proxy that is more commonly missing is the EA PGI (7.3%). This is for two reasons. First, because genetic data is not available for part of the sample (3.4%). Second, we do not calculate PGIs for those with non-European ancestries (3.9%).²¹ As a result, any analyses in this paper based on the PGI only correspond to the European-ancestry sample, though other analyses correspond to the full sample, regardless of ancestry. Family size and neighborhood SES are missing for 1.4% and 4.3% of our sample, respectively. Appendix E shows that the fractions missing each proxy are orthogonal to the schooling reform we study.

B. Proxies for Early-Life Advantages Predict Later-Life SES

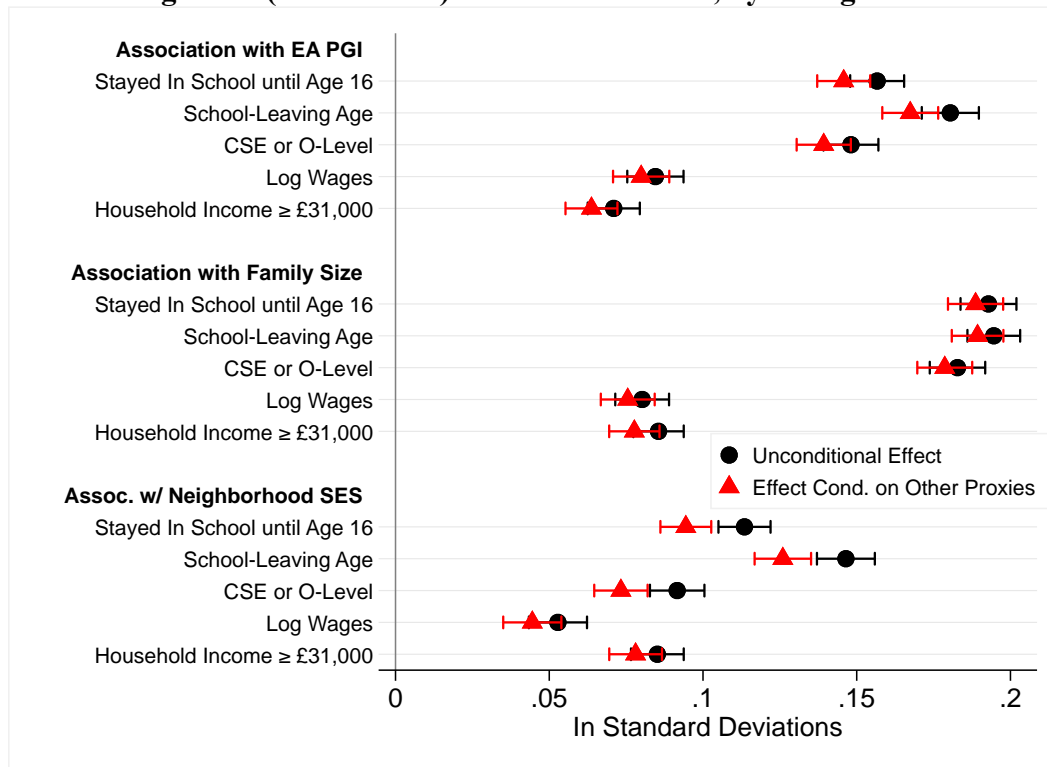
A large body of work shows that genetics, family size, and neighborhood characteristics are associated with SES in adulthood (e.g., Lee et al. 2018; Black et al. 2005; Mogstad and Wiswall 2016; Chetty and Hendren 2018a, 2018b). We begin by documenting these associations in our data. Figure 1 shows the associations of our proxies of disadvantage with different SES measures for cohorts unaffected by the reform. Both the SES measures and the proxies were standardized. The black dots show the unconditional associations. The red circles show the association of a given proxy conditional on the other two proxies. The correlation between the EA PGI and family size is 0.02; the correlation between the EA PGI and neighborhood SES is 0.07; and the correlation between family size and neighborhood SES is 0.06.

There are two main takeaways. The three proxies – each capturing a different spectrum of disadvantages at the individual, family, and neighborhood levels – are strong predictors of the different measures of long-term SES. They are also independently associated with SES: the association of a given proxy with a SES measure changes very little when we condition on the

²¹ Samples in genetic research are nearly always restricted to individuals with similar genetic ancestries—defined as a group with tightly clustered first and second principle components of their genetic data. Due to Euro-centric bias in data collection, most currently published GWASs, including Lee et al. (2018), are based on samples with “European” ancestries. These are the primary ancestries found in individuals who self-identify as “White.” As a result, PGIs based on currently available GWAS coefficients are substantially more predictive in European-ancestry samples (Martin et al. 2019) and may not generalize to groups with African, Asian, Hispanic, or other non-European ancestries. This is not a major issue in our context since over 96% of our UKB sample born between 1947 and 1967 is of European Ancestry.

other two proxies. The low correlations between the proxies indicate that they represent distinct sources of advantages.²² Notice how the EA PGI is as predictive as the other two proxies.

Figure 1: (Pre-Reform) Differences in SES, by Background



Notes: The figure shows coefficients of regressions of the dependent variables listed in the rows on the proxies for disadvantage (the regressions also include pre-reform birth cohort trends). The black circles display coefficients from regressions including just one of the three proxies. The red triangles display coefficients from regressions including all three proxies. Both the dependent variables and the proxies were normalized. The brackets show 95% confidence intervals. The sample is restricted to those born between September 1, 1947 and August 31, 1957 (i.e. the cohorts unaffected by the ROSLA). Wages are imputed based on one's occupation. Sample sizes vary between 63,172 and 78,764 depending on the measure of SES.

Again, these associations do not represent causal relationships between the proxies and later-life SES (e.g., Kong et al. 2018; Angrist et al. 2010). A separate discussion about the interpretation of the association between the EA PGI and SES is warranted.

²² To illustrate, imagine to the contrary that the EA PGI and the neighborhood SES were noisy proxies for the same latent measure of advantage and that the low correlation between the two proxies was entirely due to uncorrelated noise. In Figure 3, we estimate that the standardized association between the PGI and school leaving age is 0.31. Using a classical measurement error correction formula, this would imply that the standardized association between the (noiseless) latent measure and school-leaving age would have to be $0.31/0.07 = 4.43$. However, standardized associations cannot be greater than one unless they are strongly negatively correlated with other covariates, which is not the case in our setting.

3. Interpreting the EA PGI and Its Relationship with Later-Life SES

The interpretation of a PGI is subtle and complicated. For example, it is many people’s instinct to interpret the EA PGI as a measure of innate ability. Although a PGI can capture these types of individual attributes, the PGI will also be correlated with other individual, familial, and community characteristics. It is, therefore, incorrect to think of it solely as a measure of innate ability.

A. The Association between the EA PGI and SES

Consider a simple regression of an SES outcome on the PGI:

$$Y_i = \kappa_0 + \kappa_1 PGI_i + \chi_i. \quad (2)$$

There are at least two reasons why κ_1 will be a biased estimate of the causal effect of the PGI on SES: population stratification and indirect genetic effects. Population stratification is a form of omitted variable bias where specific genetic variants are more common in a particular group and the average of the outcome of interest Y is higher (or lower) among this group than in the rest of the population. These genetic variants will predict the outcome of interest, but the relationship is not causal.²³ Empirically, it has been shown that controlling for genetic principle components removes much but not all of the bias due to population stratification (Price et al. 2006). We show in the Robustness Section that our results are qualitatively the same when we control for the first 20 principal components of the participants’ genetic matrix.

Another omitted variable in regression (2) is parental genetics: parental genetics drive both a child’s own genetics (since the child inherits her genetic variants from her parents) and the child’s SES (through parental behaviors and characteristics, such as parental education). Since the parents’ and child’s PGIs are correlated, when the analysis does not control for parental genotypes, as in Figure 1, the effects of the parental genotypes are partly captured by the offspring’s PGI.

²³ The canonical example of population stratification is chopstick use (Price et al. 2006). Consider a GWAS of whether an individual regularly uses chopsticks and a sample that includes individuals of Asian and non-Asian backgrounds. This GWAS would find many SNPs that are associated with chopstick use, but each of these associations would correspond to SNPs that have alleles that are more common in Asian populations rather than to SNPs that represent any sort of genetic pathway between genes and chopstick use.

B. Controlling for the PGIs of Parents

By using information on the PGIs of an individual’s parents, we can however isolate variation in the individual’s PGI that is random and therefore orthogonal to her environment. This is possible due to Mendelian inheritance: a biological phenomenon that causes PGIs to be randomly assigned to children conditional on the sum of the PGIs of their parents (Kong et al. 2020). We estimate the following model:

$$Y_i = \phi_0 + \phi_1 PGI_i + \phi_2 Parental PGI_i + e_i. \quad (3)$$

where $Parental PGI_i$ is the sum of the PGIs of individual i ’s parents.

The coefficient on the PGI, ϕ_1 , estimates the “direct genetic effects.”²⁴ Direct genetic effects are defined as the causal effect of inheriting specific genetic variants holding constant one’s environment.²⁵ In contrast, ϕ_2 primarily picks up “indirect genetic effects” from one’s parents. These are the effects of genetic variants of the parents that were not inherited by the child but still affect her because they influence her parents’ characteristics and behaviors (Kong et al. 2018). Notice that we do not have exogenous variation in the parental PGI so it may also represent other environmental advantages, such as indirect genetic effects from other relatives and population stratification.

It is important to stress that the direct genetic effects do not necessarily reflect “ability” or “innate” individual characteristics that may influence productivity. Although the approach identifies the causal effects of genetic variation on SES, we cannot pin down the channel(s) of these effects. For example, suppose that height has a causal effect on wages because *ceteris paribus* employers favor taller people (Persico et al. 2004; Case and Paxson 2008). Then any genetic variant that causes an increase in height will also cause an increase in wages. EA and height have been shown to be genetically correlated (Okbay et al. 2016), such that individuals with higher EA

²⁴ It is also possible to estimate direct genetic effects using a fixed-effect model. Using a fixed-effect model, however, would require us to drop siblings with any missing data or who do not satisfy our sample restrictions. In contrast, our specification only requires genetic data for parents, which can be imputed using only genetic data from at least one other sibling or in some cases was observed because one parent was genotyped. This increases our sample size. Also, our specification allows us to directly measure the association between the parental PGI and SES and to measure the change in this association as a result of the 1972 ROSLA.

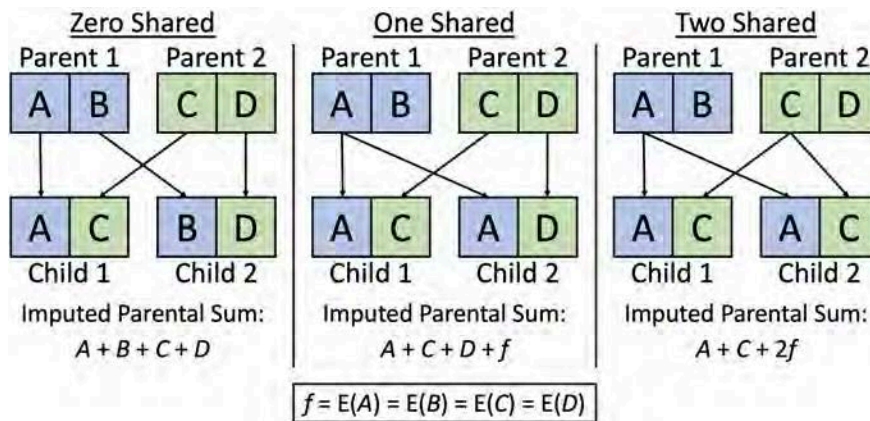
²⁵ Direct genetic effects are still a function of the environment, however; environmental factors can change the impact of specific genetic variants. Indeed, an important result of this paper is that education can change the direct effect of the EA PGI.

PGIs will tend to be taller and to earn higher wages, but, in this example, this relationship does not necessarily imply that individuals with higher EA PGIs are more productive. However, if some portion of the EA PGI represents “innate” characteristics, this portion would be captured by the direct genetic effect and not the indirect/environmental effect.

C. Imputing the PGIs of Parents

While few study participants had one of their parents genotyped, it is possible to impute the sum of the parental genotypes if at least one of the study participant’s siblings were also genotyped (Young et al. 2020). An illustration of this imputation procedure is found in Figure 2. Imagine we would like to impute the sum of the genotypes of a participant’s parents at some SNP. Because people normally have two copies of each chromosome, we represent the genotype of Parent 1 and Parent 2 at this SNP as $(A + B)$ and $(C + D)$, respectively, where A , B , C , and D are indicators of whether a parent has a copy of a specified reference allele at the SNP on a certain copy of their chromosome. The sum of the parental genotypes at this locus is therefore $(A + B + C + D)$.

Figure 2: Illustration Imputation of Parental PGIs



Note: The figure illustrates the method used to impute the PGI of parents. In the “Zero Shared” case, both siblings inherited different alleles from each parent, such that all four parental alleles are observed in the two siblings. In the second (“One Shared”) case, both siblings inherited the same allele from one parent and a different allele from the other, in which case only three parental alleles are observed. Finally, in the “Two Shared” case, only two parental alleles are observed because both siblings inherited the same allele from each parent. In the last two cases, we impute the unobserved genotypes using the mean genotype for a single allele at that SNP. This mean genotype corresponds to the frequency at which the allele is found at that SNP on a single copy of a chromosome drawn from the population.

Say we observe two (full) siblings in our data. Each child randomly inherits one of the alleles from Parent 1 (either A or B) and one allele from Parent 2 (either C or D). As Figure 2 illustrates,

there are three possible cases. In the first case (“Zero Shared”), both siblings inherited different alleles from each parent, such that all four parental alleles are observed in the two siblings. In the second (“One Shared”), both siblings inherited the same allele from one parent and a different allele from the other, in which case only three parental alleles are observed. Finally, in the third case (“Two Shared”), only two parental alleles are observed because both siblings inherited the same allele from each parent.

It is possible to determine which case we are in at a particular SNP. In the second and third cases, we take the sum of the observed genotypes and impute the unobserved one using the mean genotype for a single allele at that SNP. This mean genotype is sometimes called the *allele frequency* in the population since it corresponds to the frequency at which the allele is found at that SNP on a single copy of a chromosome drawn from the population.^{26,27}

The imputation will lead to some error in the parental genotype measure, but it is not classical measurement error. Specifically, under a classical measurement error model, it is assumed that the variable is equal to the true value plus some amount of uninformative noise. In the case of our imputed parental genotypes, our variable is equal to the true value minus some information—the unobserved parental genotypes. Importantly, if the parental PGI is omitted, the portion of the parental PGI that biases the coefficients of (2) are the shared genetic variants between the parents and children. This is the portion of the PGI that we are able to impute without error. This means that, by controlling for the imputed parental PGIs, we can estimate the direct effect, ϕ_1 , without bias. Furthermore, since all of the variation in the imputed parental PGI corresponds to actual variation in the true parental PGI (as opposed to corresponding to uninformative noise), we can likewise estimate the indirect genetic effect without bias. More details and proofs of these properties can be found in Young et al. (2020).


²⁶ We follow a similar procedure when one parent is genotyped but the other parent is unobserved. In these cases, we always observe three of the four parental alleles, so the imputed parental sum is the sum of the three observed genotypes plus a constant equal to the frequency of the missing genotype.

²⁷ This imputation of the unobserved allele is justified under an assumption that the observed alleles are uncorrelated with the unobserved alleles. This assumption would hold if there were random mating in the population, which has been shown to not be the case. However, it has been shown that the expected correlation between parental genotypes would be on the order of 10^{-6} if parents were sorting at rates implied by the most extreme estimates. We therefore believe that the zero-correlation assumption is a reasonable approximation in this case.

D. Direct Genetic Effects

The red triangles in Figure 3 show estimates of direct genetic effects. The black circles display associations between the EA PGI and the outcomes of interest. The direct genetic effects are obtained by estimating equation (3) and plotting $\hat{\phi}_1$. The associations correspond to estimating equation (2), which does not control for the parental PGI. For comparison, all outcome variables were standardized. The PGI was also standardized. To maximize statistical power, this particular analysis uses all 33,277 UKB participants of European ancestries with a parent or sibling who had also participated in the UKB.²⁸

Figure 3: Direct Genetic Effects vs. EA PGI Associations

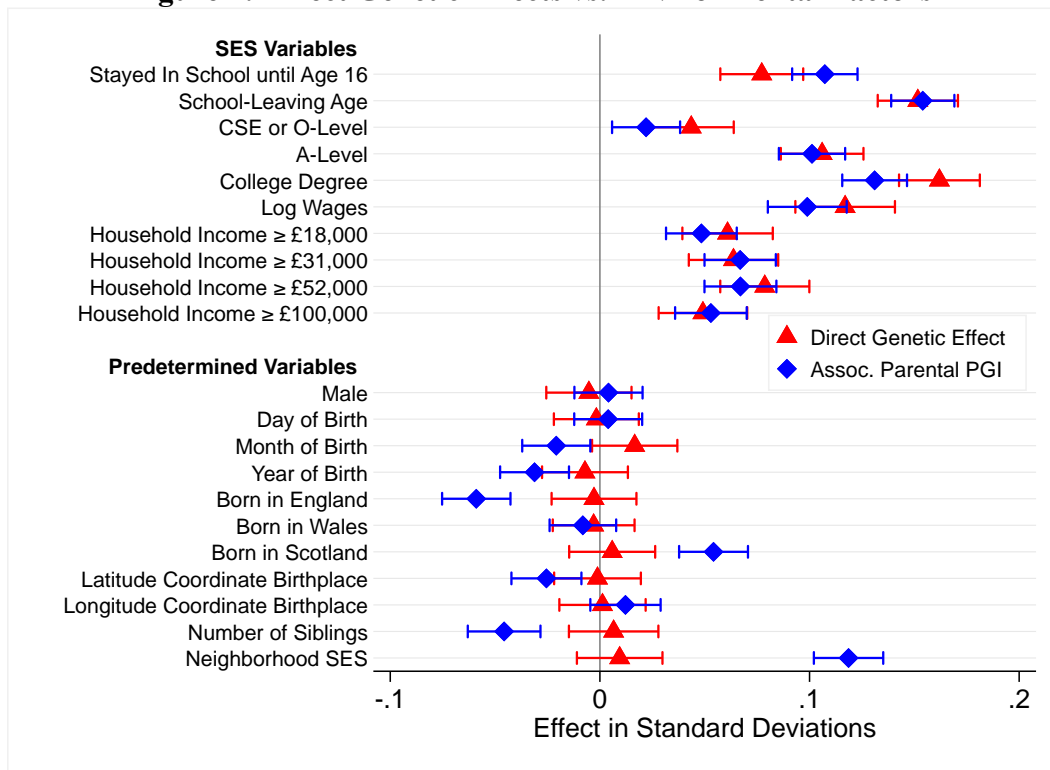


Notes: The red triangles shows estimates of the causal effect of the EA PGI. The black circles show associations with the EA PGI. The brackets show 95% confidence intervals. The EA PGI, the SES variables, and the predetermined variables were all normalized. Wages are imputed based on one's occupation. The sample includes all UKB participants of European ancestries for whom the parental EA PGI was available. The number of observations ranges from 23,298 to 33,277 depending on the outcome.

²⁸ Figures 1 and 3 were constructed using different samples (e.g., Figure 1's sample includes only those born in England, Scotland, or Wales between September 1, 1947 and August 31, 1957 and who dropped out before age 19 while Figure 3 includes only study UK Biobank participants for whom parents' genotypes were available or could be imputed). That is why the associations between the EA PGI and the SES outcomes (black circles on both figures) are slightly different in the two figures.

The top panel of the figure shows that, not only the EA PGI predicts several SES outcomes, but that it also has a causal effect on all of them. An increase of one SD in the EA PGI *causes* an increase of 0.15 SD in school-leaving age and an increase of 0.12 SD in log wages. In comparison, an increase of one SD in the EA PGI is *associated* with a 0.31 SD increase in school-leaving age and with a 0.22 SD increase in log wages. The bottom panel of the figure reports results for variables that would have been fixed before or very soon after the genotype, and consequently the PGI, was fixed. Therefore, one would expect no causal relationship between the EA PGI and these variables. Even though the EA PGI is (unconditionally) associated with several of these variables, we cannot reject that the direct genetic effect on any of them is different from zero. This can be seen as a “balance test” that the parental PGI imputation described in Section 3.C jointly with Mendelian inheritance worked as intended and that the conditional EA PGI is indeed random.

Figure 4: Direct Genetic Effects vs. Environmental Factors



Notes: The red triangles shows estimates of the causal effect of the EA PGI. The blue diamonds show associations with the parental EA PGI. The brackets show 95% confidence intervals. The EA PGI, the parental EA PGI, the SES variables, and the predetermined variables were all normalized. Wages are imputed based on one’s occupation. The sample includes all UKB participants of European ancestries for whom the parental EA PGI was available. The number of observations ranges from 23,298 to 33,277 depending on the outcome.

The associations with the SES outcomes are larger than the direct genetic effects partly because of indirect genetic effects. The blue diamonds in Figure 4 display the coefficients on the parental PGI when estimating equation (3), i.e., $\hat{\phi}_2$. Having parents with higher EA PGIs is associated with achieving higher success in adulthood (top panel). It is also associated with having fewer siblings and with living in a higher-SES neighborhood (bottom panel). Because we do not have exogenous variation in the parental PGI, these associations may also represent other environmental advantages, such as indirect genetic effects from other relatives and population stratification.

The direct genetic effects from Figure 3 are reproduced in Figure 4 for purposes of comparison. The results *suggest* the indirect genetic effects may be as large as the direct genetic effects, which is broadly consistent with previous findings (Kong et al. 2018; Howe et al. 2021; Lee et al. 2018; Cheesman et al. 2020). We will use the same design in Section 4.D to distinguish between direct genetic effects and indirect genetic effects and other environmental influences.

4. Education and the Relationship between Early-Life Advantages and Later-Life SES

A. The 1972 Raising of the School-Leaving Age

If, as believed, education levels the playing field, then we would expect it to reduce the gaps between children from disadvantaged and advantaged backgrounds. To investigate this hypothesis, we study a well-known compulsory schooling reform and its effects on these disparities. The 1972 Raising of School Leaving Age (ROSLA) legislation increased the minimum school-leaving age in England, Scotland, and Wales from 15 to 16 years of age.²⁹ These laws and their implementation have been extensively documented before (see Clark and Royer 2010, 2013; Grenet 2013; Davies et al. 2018) so we only include a brief summary of its main features here. The change took effect in September 1, 1972, implying that those who were 15 or younger before that date (i.e., born on September 1, 1957 or later) had to stay in school until at least age 16 (hereafter, we use the term “stayed in school until age 16” to refer to those who stayed in school until *at least* age 16). Infrastructure investments, such as school building to absorb the additional students, preceded the 1972 ROSLA but key elements of the school system did not change with the reform.

²⁹ http://www.legislation.gov.uk/uksi/1972/444/pdfs/uksi_19720444_en.pdf
http://www.legislation.gov.uk/uksi/1972/59/pdfs/uksi_19720059_en.pdf

Politically, there were a number of different justifications for the 1972 ROSLA including addressing the British economy’s failing competitiveness, raising England’s level of education to be on par with other countries, and reducing the number of young people seeking employment at high unemployment times. Interestingly, a 1959 report by the ministry of education advising the Government to increase compulsory schooling to age 16 highlighted the need to tap “all the available supply of talent” and the social barriers to doing so. The Crowther Report documented that “half of the National Service recruits to the Army who were rated in the two highest ability groups had left school at age 15” and that among recruits “coming from families of manual workers two-thirds of those in the two highest ability groups had left school at 15.”³⁰

In our analysis, we restrict the sample to participants born between September 1, 1947 and August 31, 1967 who were either born in England, Scotland, or Wales or immigrated to the UK before age 15. We also restrict the sample to individuals who left school at age 18 or younger to increase the precision of our estimates (Banks and Mazzonna 2012; Banks et al. 2019). We confirm previous results (e.g., Clark and Royer 2013) that the ROSLA did not induce students to stay in school past age 18 (see Appendix L), such that this sample restriction does not bias our estimates. Indeed, we show in the Robustness Section that that the results are qualitative the same if this sample restriction is relaxed.

Figure 5 illustrates some of the impacts of the ROSLA. It shows the fraction of study participants who stayed in school until age 16 (black circles); age 17 (red Xs); and age 18 (blue triangles) by year of birth. Year of birth runs from September 1 of a given year to August 31 of the following year. For example, those born in year 0 correspond to those born between September 1, 1957 and August 31, 1958, the first cohort affected by the ROSLA. The figure shows linear trends for a 4-year bandwidth around September 1, 1957 and quadratic dates for a 10-year bandwidth.


The figure shows that the ROSLA generated a discontinuity in the relationship between staying in school until age 16 and date of birth. There is a sharp increase for those born after September 1, 1957. The discontinuities at ages 17 and age 18 are much smaller but still significant.

We estimate the ROSLA increased the fraction of those staying in school until age 16 by 25 percentage points—see Appendix Table B1. The fraction staying until ages 17 and 18 increased

³⁰ See <http://www.educationengland.org.uk/documents/crowther/crowther1959-1.html> pages 131-132

by 3-4 and 2-3 percentage points. On average, the cohort affected by the ROSLA has approximately a third of a year of schooling more than those who could drop out at age 15.³¹ This estimate is larger than estimates reported in other studies considering this same reform because we restrict our data to those who report leaving school at or before age 18.

Figure 5: Educational Attainment by Year of Birth



Notes: The figure shows the fraction of study participants who stayed in school until age 16 (black circles), age 17 (red crosses), and age 18 (blue triangles) by year of birth. Year of birth runs from September 1 of a given a year to August 31 of the following year. Those born in Year 0 were born between September 1, 1957 and August 31, 1958. Cohorts born after Year 0 had to stay in school until age 16 while cohorts born before could leave at age 15. The figure shows linear trends for a 4-year bandwidth and quadratic trends for a 10-year bandwidth. We use triangular kernel weights that give greater weight to study participants born closer to September 1, 1957. $N = 46,308$ (4-year bandwidth) and 114,025 (10-year bandwidth).

The increase in schooling generated by the ROSLA resulted in an increase in wages and income. Appendix Figure B1 shows that those born after September 1, 1957, who, as a result of the 1972 ROSLA, ended up with more schooling than their counterparts, earn higher wages and have higher incomes. We estimate that an additional year of secondary education increased

³¹ Importantly, the F-stat on the effect of the ROSLA on school-leaving age is greater than 200 in all of specifications, implying that our subsequent two-stage least squares estimates will be well-powered (Lee et al., 2020).

middle-age wages by 5% – see Appendix Table B2. These are consistent with Grenet (2013) who estimated a return of 6-7%.^{32,33}

Appendix A shows that predetermined characteristics and our proxies for early-life advantages are smooth around the September 1, 1957 cutoff. These results suggest that the identifying assumption of the regression discontinuity design is satisfied and that we can stratify the results by our advantage proxies.

B. Did the 1972 Raising of the School-Leaving Age Reduce Educational Gaps?

One of the rationales for the 1972 ROSLA was to reduce educational gaps between children from different socioeconomic backgrounds and to allow high-ability, disadvantaged students to fulfill their potential (Crowther Report 1959). To investigate whether the ROSLA was effective in this regard, Figure 6 shows the average school-leaving age by year of birth for those in the top and bottom terciles of the distributions of the EA PGI (top panel), family size (middle panel), and neighborhood SES (bottom panel). Appendix D shows comparable figures for the fraction of students who stayed in school until ages 16, 17, and 18 – Appendix Figure D1 is of particular interest because age 16 was the margin targeted by the reform.

Figure 6 indicates that the ROSLA was effective in reducing the educational disparities between children from less and more advantaged backgrounds. While the average school-leaving age of both the bottom and top terciles jump discontinuously after September 1, 1957, the jump is larger for the bottom tercile (i.e., the children from more disadvantaged backgrounds), such that the difference between the red and black lines diminishes noticeably. This is true for all three different proxies of disadvantages: individual, family, or neighborhood. The gap in school-leaving age decreased from 0.44 to 0.33 of a year of schooling between those in the top and the bottom terciles of the EA PGI distribution; from 0.44 to 0.25 of a year of schooling between those in the top and the bottom terciles of the family size distribution; and from 0.35 to 0.24 of a year of schooling between those in the top and the bottom terciles of the neighborhood SES distribution. These are large reductions considering that the ROSLA increased school-leaving age on average by a third of a year of schooling—see Appendix Table B1.

³² Grenet (2013) studied the 1972 ROSLA using data from the Quarterly Labour Force Survey.

³³ The effects are robust to the alternative bandwidths. Interestingly, these 2SLS estimates are also consistent with the OLS estimates (i.e., regressing log-wages directly onto the endogenous SLA variable), suggesting a limited role for omitted variable bias in this context.

The trends shown in Figure 6 are obtained by estimating the following regressions:

$$SLA_i = \delta_0 + \delta_1 Post_i + \mathbf{B}'_i \boldsymbol{\delta}_2 + Post_i \times \mathbf{B}'_i \boldsymbol{\delta}_3 + k(DoB_i) + \varepsilon_i. \quad (4)$$

where SLA_i is individual i 's school-leaving age; $Post_i$ is 1 if individual i was born on or after September 1, 1957 (and 0 otherwise); \mathbf{B} is a vector containing proxies for early-life advantages; and DoB_i is individual i 's date of birth. Date of birth is measured in days relative to the cutoff, such that $DoB = 0$ for someone born on September 1, 1957. The function $k(\cdot)$ captures birth cohort trends in educational attainment, which are allowed to differ on either side of the September 1, 1957 cutoff. We run two alternative specifications: (1) a 10-year bandwidth with quadratic trends in date of birth and (2) a 4-year bandwidth with linear trends.³⁴ We use triangular kernel weights that give greater weight to study participants born closer to the cutoff.³⁵

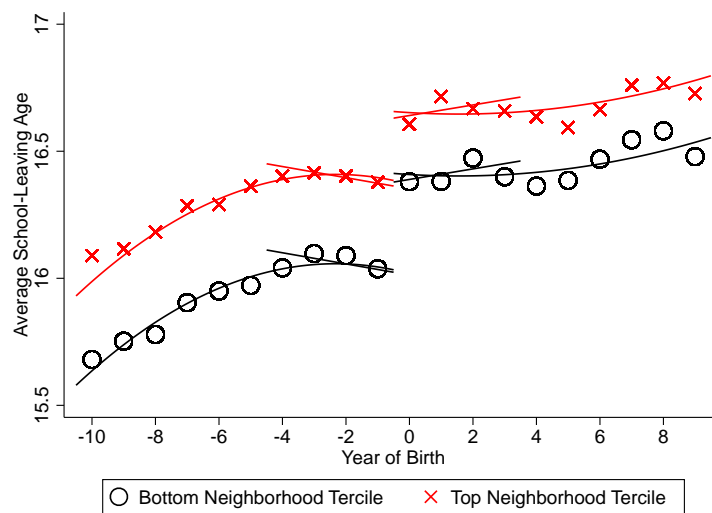
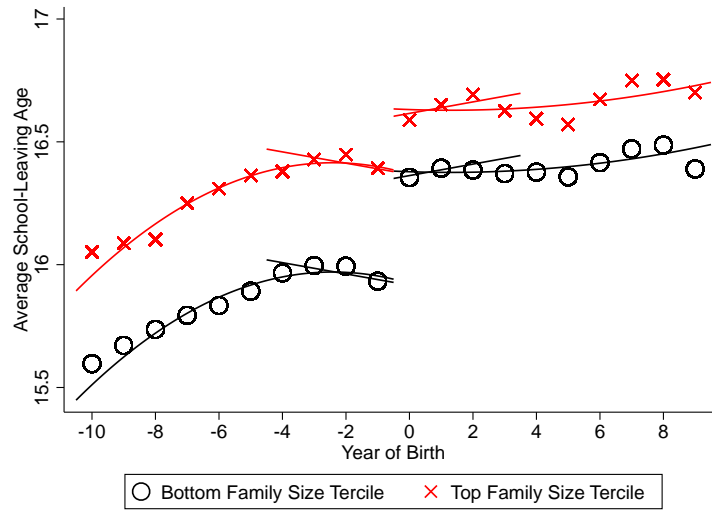
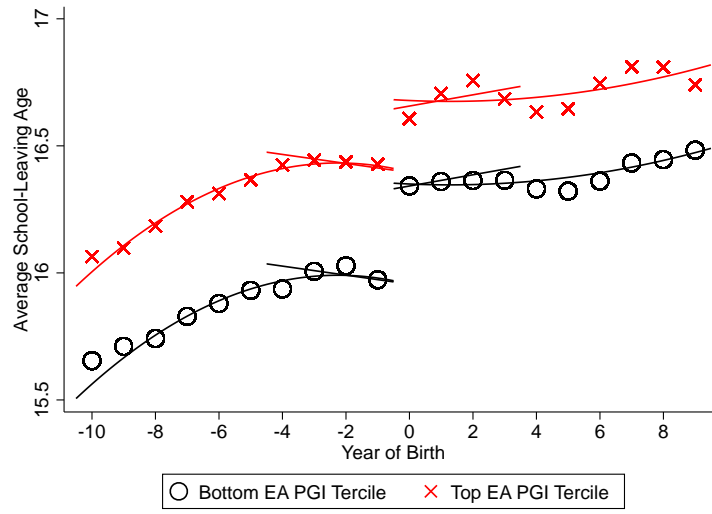
We use a discrete specification in Figure 6, where \mathbf{B}_i is a vector containing indicators for being in the top and in the middle terciles of the distributions of EA PGI (top panel), family size (middle panel), and neighborhood SES (bottom panel). Notice that the birth cohort trends $k(DoB_i)$ are assumed to be the same for individuals in the three different terciles of a given proxy distribution.³⁶ The Robustness Section discusses that – although relaxing this assumption implies a large loss in precision – the point estimates results are overall very similar. Middle-tercile participants were also used to estimate Figure 6's birth cohort trends. They are not shown in the figure for ease of exposition; this is true not only for Figure 6 but for all figures comparing the bottom and top terciles.

³⁴ To the best of our knowledge, methods for optimal bandwidth selection for interaction effects have yet to be developed. We worry that the standard CCT optimal bandwidth is too conservative in the case of interactions which is why we consider both a 4-year (the CCT optimal bandwidth) and a 10-year bandwidth. In the Robustness Section we show that our results are robust to the choice of bandwidth.

³⁵ Even though previous work studying the 1972 ROSLA clustered standard errors by month-year of birth (e.g., Clark and Royer 2013; Davies et al. 2018), we do not need to cluster our standard errors because our data include exact date of birth (Lee and Card 2008).

³⁶ Notably, this specification excludes any additional control variables, such as principal components of the genetic data. Principal components are normally included as covariates in studies that include PGIs to absorb variation in the PGI and outcome variables related to ancestry differences. Since we do not include principal components in our models, the PGI may also capture at-birth advantages related to ancestry. In the Robustness Section, we show that including principal components has only a negligible effect on our results.

Figure 6: Effect of the 1972 ROSLA on School-Leaving Age, by EA PGI, Family Size, and Neighborhood SES



Notes: The figures show average school-leaving age by year of birth, separately for those in the bottom tercile (black circles) and in the top tercile (red Xs) of the following distributions: EA PGI (top panel); family size (middle panel); and neighborhood SES (bottom panel). Year of birth runs from September 1 of a given year to August 31 of the following year. Those born in Year 0 were born between September 1, 1957 and August 31, 1958. Cohorts born after Year 0 had to stay in school until age 16 while cohorts born before could leave at age 15. The figure shows linear cohort trends for a 4-year bandwidth and quadratic cohort trends for a 10-year bandwidth. For a given proxy, all three terciles share the same cohort trends. We use triangular kernel weights that give greater weight to study participants born closer to September 1, 1957. $N = 42,731$ (EA PGI with 4-year bandwidth); 105,693 (EA PGI with 10-year bandwidth); 45,701 (family size with 4-year bandwidth); 112,395 (family size with 10-year bandwidth); 44,043 (neighborhood SES with 4-year bandwidth); and 109,177 (neighborhood SES with 10-year bandwidth).

Table 1: Effect of the 1972 ROSLA on School-Leaving Age, by EA PGI, Family Size, and Neighborhood SES

Top Panel: Discrete	School-Leaving Age							
	4-Year Bandwidth				10-Year Bandwidth			
Top PGI * Post	-0.125			-0.118	-0.113			-0.101
	[0.027]			[0.026]	[0.017]			[0.017]
Top Family * Post		-0.198		-0.197		-0.192		-0.185
		[0.025]		[0.024]		[0.016]		[0.016]
Top Neighborhood * Post			-0.088	-0.075			-0.108	-0.096
			[0.028]	[0.027]			[0.017]	[0.017]
Mid PGI * Post	-0.065			-0.059	-0.060			-0.053
	[0.027]			[0.026]	[0.017]			[0.017]
Mid Family * Post		-0.157		-0.156		-0.147		-0.138
		[0.027]		[0.027]		[0.018]		[0.017]
Mid Neighborhood * Post			-0.037	-0.042			-0.031	-0.034
			[0.026]	[0.025]			[0.017]	[0.016]
Post	0.368	0.421	0.353	0.521	0.386	0.438	0.379	0.530
	[0.025]	[0.023]	[0.025]	[0.030]	[0.021]	[0.020]	[0.021]	[0.024]
Bottom Panel: Continuous								
PGI * Post	-0.048			-0.045	-0.048			-0.042
	[0.011]			[0.011]	[0.007]			[0.007]
Family * Post		-0.088		-0.088		-0.082		-0.080
		[0.010]		[0.010]		[0.007]		[0.007]
Neighborhood * Post			-0.040	-0.033			-0.049	-0.043
			[0.011]	[0.011]			[0.007]	[0.007]
Post	0.306	0.305	0.311	0.311	0.328	0.328	0.333	0.331
	[0.020]	[0.020]	[0.020]	[0.020]	[0.019]	[0.019]	[0.019]	[0.018]
Observations	46,308	46,308	46,308	46,308	114,025	114,025	114,025	114,025

Notes: The dependent variable is school-leaving age. PGI is the polygenic index for educational attainment. Post is an indicator for being born on or after September 1, 1957. In the top panel, Top and Mid are indicators for being in the top or middle terciles of the distribution of a given proxy. In the bottom panel, PGI, Family, and Neighborhood are all standardized to have mean zero and standard deviation of 1. The first four columns include linear trends in exact date of birth while the last four include quadratic trends. In both cases, trends are allowed to be different before and after September 1, 1957. Robust standard errors between brackets. In the top panel, we omit the coefficients on Top PGI, Top Family, Top Neighborhood, Mid PGI, Mid Family, and Mid Neighborhood. In the bottom panel, we omit the coefficients on PGI, Family, and Neighborhood.

The results from this discrete specification are shown in the top panel of Table 1. In the bottom panel of Table 1 we report results from a continuous specification of equation (4), where \mathbf{B}_i is a vector containing continuous (standardized) measures of the EA PGI, family size, and/or neighborhood SES. Because of space constraints, the estimated coefficients δ_0 and δ_2 are not reported in the table.

The results in Table 1 confirm that the 1972 ROSLA was somewhat successful at leveling the playing field, at least in terms of schooling. The interactions of the indicator for being born after September 1, 1957 with the EA PGI, family size, and with neighborhood SES are all negative with most being statistically significant at the 1%.

These reductions in educational differences may not be very meaningful if the students forced to stay in school an extra year put in little effort and gained little with the additional schooling. However, one potential concrete benefit of this additional year of education was the opportunity to sit for a set of qualification exams that are taken at the end of grade ten (when participants are typically 16): the Certificate of Secondary Education (CSE) or the General Certificate of Education (GCE) Ordinary Level (also known as an O-Level). By compelling students to stay in school, the 1972 ROSLA led students to complete these courses and get these qualifications, which are valued in the labor market (Dickson and Smith 2011). Appendix Figure D4 shows that the 1972 ROSLA had a larger effect on those in the bottom terciles of the three distributions, reducing the disparities in qualifications as a consequence. This is consistent with the estimated reduction in school-leaving age disparities and provides suggestive evidence that the students compelled to stay in school for an additional year acquired concrete benefits as a result.

C. Did the 1972 Raising of the School-Leaving Age Reduce Wage Gaps?

We documented in the previous section that the ROSLA was successful in reducing educational disparities between children from different backgrounds. What would be the expected impact of such reduction on wage disparities *if the returns to schooling were the same for everyone?* At the eve of the program, the wage gap between the top and bottom terciles of each proxy distribution was 0.066 (EA PGI), 0.055 (family size), and 0.040 log points (neighborhood SES). Table 1 estimated the ROSLA reduced the top-bottom gap in school-leaving age by 0.113 (EA PGI), 0.192 (family size), and 0.108 years of schooling (neighborhood SES). If the return to these additional years of schooling were 5% (see Appendix Table B2), we would expect the wage

gaps described above to decrease by 8.6% (EA PGI),³⁷ 17.5% (family size), and 13.5% (neighborhood SES).

In practice, we find no evidence that there was a reduction in the wage disparities. On the contrary, if anything, there was a widening of such gaps. Figure 7 shows the average log wage by year of birth, separately for those in the bottom and in the top terciles of each proxy distribution. The ROSLA increased wages for both the bottom and top terciles, but, different than expected, the increase was not larger for those in the bottom terciles of the three distributions. For the EA PGI in particular, the increase for the bottom tercile is visually smaller than the increase for the top tercile. Appendix Table F1 confirms these results.

Figure 8 summarizes the first-stage (Table 1) and reduced-form results (Appendix Table F1). For any given proxy – EA PGI (black), family size (blue), or neighborhood SES (red), it reports the difference between those in the top and bottom terciles of the proxy distribution on the eve of the ROSLA (Xs) and immediately after (circles). The left y-axis and left panel show top-bottom differences in the schooling-leaving age. The right y-axis and right panel show top-bottom differences in log wages. The brackets report 95% confidence intervals testing whether the post-reform difference is equal to the pre-reform difference.³⁸

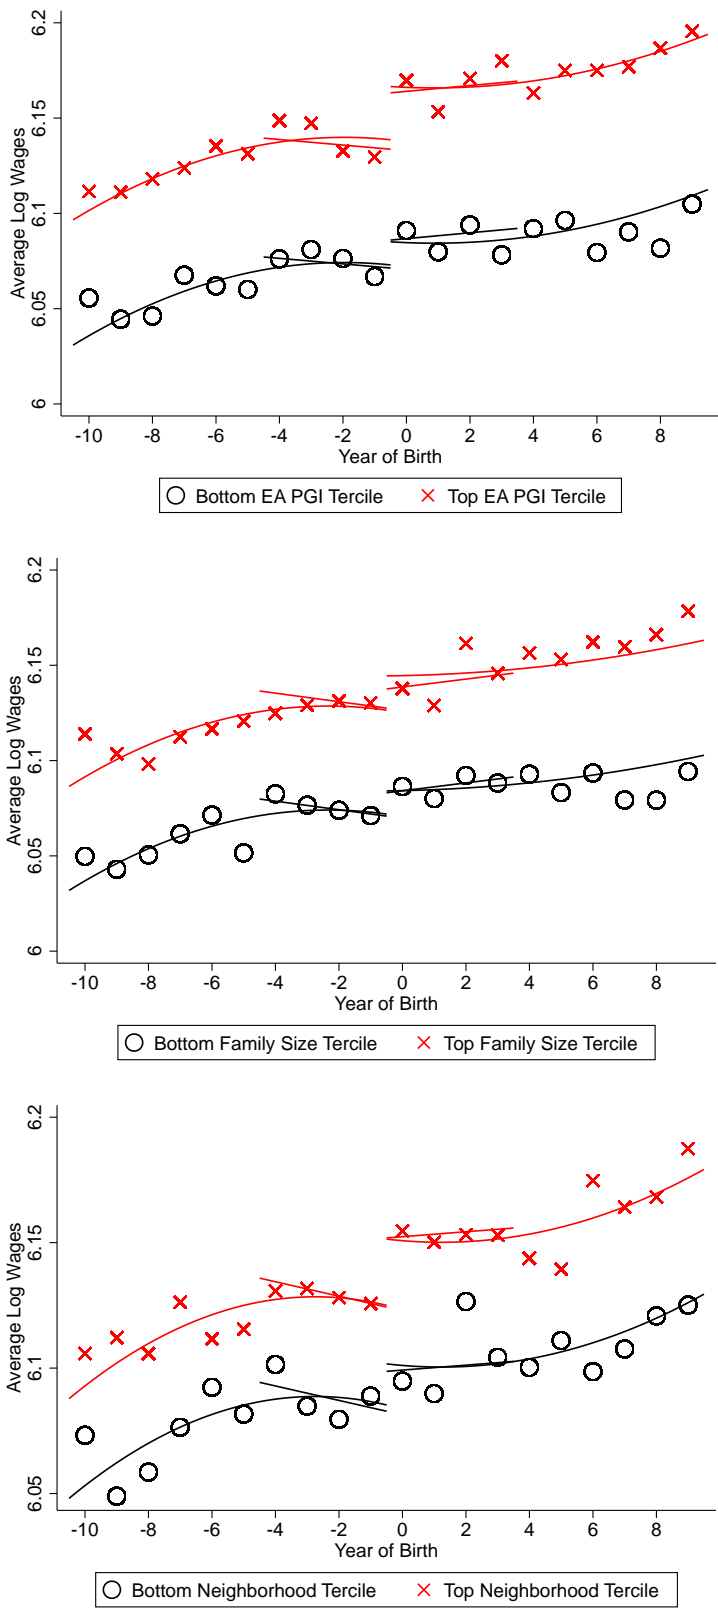
The left panel shows that the 1972 ROSLA was effective in reducing disparities in education. The pre-ROSLA top-bottom differences in school-leaving age decreased by 26-43% across our three proxies. The right panel shows that, contrary to what was expected, disparities in wages did not reduce. The estimates suggest instead that the wage gap did not change and may have even increased.

Our wage results are consistent with advantaged children having higher returns to their additional schooling. The technology of skill formation framework (Heckman 2000; Cunha et al. 2006) provides a rationale for why the returns to schooling may vary with the EA PGI, family size, or neighborhood SES. Conti and Heckman (2010), for example, propose that genetics may affect the stock of skills both as an input and by shaping the production function, which in turn may increase the returns to schooling (Heckman, Stixrud, and Urzua 2006; Urzua 2006).

³⁷ $0.086=(0.05)*(0.113)/0.066$, where 0.05 is the return to schooling, 0.113 is the reduction in the educational gap, and 0.066 is the original top-bottom difference by EA PGI.

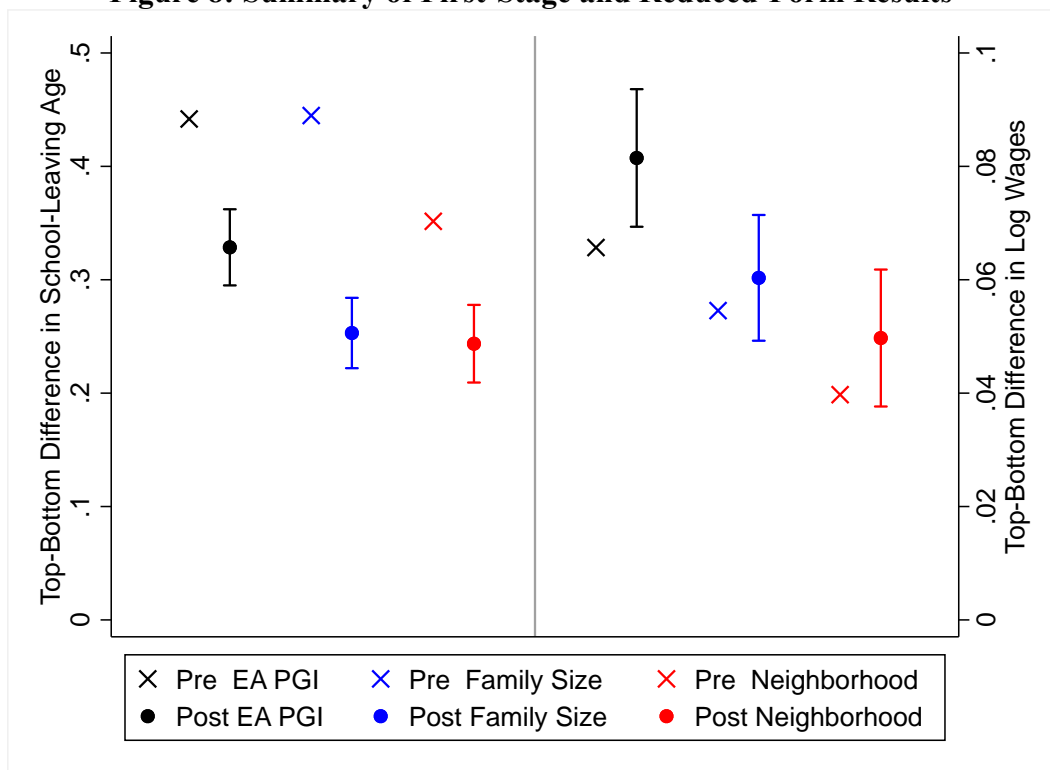
³⁸ Because the figure illustrates a test of the difference between the points indicated by the X's and the solid circles, we do not report confidence intervals for the pre-ROSLA levels.

Figure 7: Effect of the 1972 ROSLA on Log Wages, by EA PGI, Family Size, and Neighborhood SES



Notes: The figures show average log wages by year of birth, separately for those in the bottom tercile (black circles) and in the top tercile (red crosses) of the following distributions: EA PGI (top panel); family size (middle panel); and neighborhood SES (bottom panel). Wages are imputed based on one's occupation. Year of birth runs from September 1 of a given year to August 31 of the following year. Those born in Year 0 were born between September 1, 1957 and August 31, 1958. Cohorts born after Year 0 had to stay in school until age 16 while cohorts born before could leave at age 15. The figure shows linear cohort trends for a 4-year bandwidth and quadratic cohort trends for a 10-year bandwidth. For a given proxy, all three terciles share the same cohort trends. We use triangular kernel weights that give greater weight to study participants born closer to September 1, 1957. $N = 42,731$ (EA PGI with 4-year bandwidth); 105,693 (EA PGI with 10-year bandwidth); 45,701 (family size with 4-year bandwidth); 112,395 (family size with 10-year bandwidth); 44,043 (neighborhood SES with 4-year bandwidth); and 109,177 (neighborhood SES with 10-year bandwidth).

Figure 8: Summary of First-Stage and Reduced-Form Results



Notes: The figure shows differences in average school-leaving age (left y-axis) and in average log wages (right y-axis) between the top and bottom terciles of the distributions of the EA PGI (black), family size (blue), and neighborhood SES (red) at the eve of the ROSLA (Xs) and immediately after (circles). Wages are imputed based on one's occupation. The brackets report 95% confidence intervals testing whether the post-reform difference is equal to the pre-reform difference. Estimates from specification with 10-year bandwidth and quadratic trends. $N = 105,693$ (EA PGI); 112,395 (family size); 109,177 (neighborhood SES).

To estimate how the returns to schooling vary with our proxies, we estimate the following equation:

$$\ln W_i = \eta_0 + \eta_1 SLA_i + \mathbf{B}'_i \eta_2 + SLA_i \times \mathbf{B}'_i \eta_3 + l(DoB_i) + \epsilon_i, \quad (5)$$

We estimate (5) through two-stage least squares (2SLS), using the indicator for being born on or after September 1, 1957 and its interactions with \mathbf{B}_i to instrument for school-leaving age and its interactions with \mathbf{B}_i . The results are shown in Table 2.

**Table 2: Effect of an Additional Year of Schooling on Log Wages,
by EA PGI, Family Size, and Neighborhood SES**

Top Panel: Discrete	Log Wages							
	4-Year Bandwidth				10-Year Bandwidth			
Top PGI * SLA	0.082			0.083	0.061			0.061
	[0.035]			[0.036]	[0.018]			[0.018]
Top Family * SLA		0.018		0.018		0.040		0.039
		[0.034]		[0.035]		[0.019]		[0.019]
Top Neighborhood * SLA			0.057	0.051			0.044	0.041
			[0.033]	[0.033]			[0.018]	[0.018]
Mid PGI * SLA	-0.019			-0.027	0.009			0.007
	[0.028]			[0.029]	[0.015]			[0.015]
Mid Family * SLA		0.076		0.078		0.032		0.030
		[0.034]		[0.035]		[0.018]		[0.018]
Mid Neighborhood * SLA			-0.002	-0.011			0.001	0.000
			[0.028]	[0.028]			[0.015]	[0.015]
SLA	0.043	0.031	0.040	0.020	0.033	0.031	0.037	0.005
	[0.024]	[0.020]	[0.025]	[0.026]	[0.019]	[0.017]	[0.019]	[0.017]
Bottom Panel: Continuous								
PGI * SLA	0.022			0.020	0.023			0.022
	[0.012]			[0.012]	[0.007]			[0.007]
Family * SLA		0.023		0.020		0.021		0.019
		[0.011]		[0.011]		[0.006]		[0.006]
Neighborhood * SLA			0.026	0.023			0.021	0.019
			[0.014]	[0.015]			[0.008]	[0.008]
SLA	0.059	0.058	0.059	0.071	0.055	0.053	0.053	0.061
	[0.024]	[0.025]	[0.024]	[0.027]	[0.020]	[0.021]	[0.020]	[0.022]
Observations	46,308	46,308	46,308	46,308	114,025	114,025	114,025	114,025

Notes: The dependent variable is log wages, which are imputed based on one's occupation. SLA is school-leaving age. PGI is the polygenic index for educational attainment. In the top panel, Top and Mid are indicators for being in the top or middle terciles of the distribution of a given proxy. In the bottom panel, PGI, Family, and Neighborhood are all standardized to have mean zero and standard deviation of 1. The first four columns include linear trends in exact date of birth while the last four include quadratic trends. In both cases, trends are allowed to be different before and after September 1, 1957. Robust standard errors between brackets. In the top panel, we omit the coefficients on Top PGI, Top Family, Top Neighborhood, Mid PGI, Mid Family, and Mid Neighborhood. In the bottom panel, we omit the coefficients on PGI, Family, and Neighborhood.

We find larger returns to schooling for advantaged children, as measured by our proxies. The return to an additional year of schooling for those in the top terciles of the EA PGI, family size and neighborhood SES distributions was 6, 4, and 4 percentage points higher than the returns for those in the bottom terciles of these distributions (top right panel of Table 2).³⁹ Overall, an

³⁹ Appendix Tables D2 and D3 suggest that the fraction of students who stayed in school until age 17 or until age 18 because of the ROSLA may have been higher in the top tercile of the EA PGI distribution than in the bottom tercile. If the returns to staying in school until age 17 or until age 18 were higher than the return to staying until age 16, this

improvement in one standard deviation unit in any of our proxies increases the returns to schooling by approximately 2 percentage points (bottom panel of Table 2). Furthermore, the coefficients on the interactions are very similar in the specifications that include all proxy variables at once (columns 4 and 8). This suggests that the gradients in returns to schooling with respect to the three different proxies are independently important. We show in the Robustness Section that we find qualitatively similar results if we use income instead of wages as the dependent variable.

D. Genetics, Environment, and the Returns to Schooling

The results presented so far raise the question of *why* children from advantageous backgrounds benefited more from the additional schooling. For example, did the compulsory schooling unlock the individual, untapped potential of children who would have otherwise dropped out at age 15? Did it make it easier for children from advantaged families to leverage their resources and connections to get further ahead? Some would consider the latter to represent unequal opportunities but not the former. While many channels beyond these two examples may underlie the gradient in returns to schooling, we can use the random process of genetic inheritance to test whether the gradient is caused by genetically-influenced individual characteristics, including innate ability or skills.

Econometric Model

The effects of the ROSLA on the relationship between the PGI and wages were estimated in Section 4.B using:

$$Y_i = \delta_0 + \delta_1 Post_i + \delta_2 PGI_i + \delta_3 (PGI_i \times Post_i) + k(DoB_i) + \varepsilon_i. \quad (4')$$

As discussed in Section 3.B, δ_2 and δ_3 capture not only direct genetic effects but also indirect genetic effects and other environmental influences.

We can however isolate the random variation in one's PGI by conditioning on the (sum of the) PGIs of the individual's parents:

could in principle explain why the LATE estimate of the return to an additional year of schooling is higher for the top tercile of the EA PGI distribution than for the bottom tercile (see Table 2). However, OLS estimates suggest that the returns to staying in school until age 17 or until age 18 are actually lower than the return to staying until age 16 (see Appendix Table B3). While these are non-causal estimates, Appendix Table B2 shows that the OLS and 2SLS estimates of the return to an additional year of schooling are comparable.

$$\begin{aligned}
Y_i = & \theta_0 + \theta_1 Post_i + \theta_2 PGI_i + \theta_3 (PGI_i \times Post_i) + \\
& + \theta_4 Parental PGI_i + \theta_5 (Parental PGI_i \times Post_i) + m(DoB_i) + \xi_i.
\end{aligned} \tag{6}$$

The coefficient on the PGI, θ_2 , estimates the direct genetic effects before the ROSLA. In contrast, θ_4 primarily picks up pre-reform “indirect genetic effects” from one’s parents. As stated before, we do not have exogenous variation in the parental PGI so it may also represent other environmental advantages, such as indirect genetic effects from other relatives and population stratification. The coefficient on $PGI_i \times Post_i$, θ_3 , represents how the direct genetic effects changed with the ROSLA. Similarly, the coefficient on $Parental PGI_i \times Post_i$, θ_5 , represents the change in the indirect genetic and other environmental effects caused by the ROSLA.

As discussed in Section 3.C, while few study participants had one of their parents genotyped, it is possible to impute the parental genotypes if at least one of the study participant’s siblings were also genotyped (Young et al. 2020). Measured or imputed data on parental PGIs are available for about 10% of our sample.

However, we show in Appendix I how it is possible to gain statistical power by using our entire sample, including those participants for whom we do not have parental PGIs, to estimate the parameters of equation (6). In particular, we estimate the following slightly modified model:

$$\begin{aligned}
Y_i = & \theta_0 + \theta_1 Post_i + \theta_2 PGI_i + \theta_3 (PGI_i \times Post_i) + \\
& + \theta_4 [(1 - M_i) Parental PGI_i + M_i PGI_i] + \theta_5 [(1 - M_i) Parental PGI_i + M_i PGI_i] \times Post_i + \\
& + \theta_6 M_i + \theta_7 (M_i \times Post_i) + m(DoB_i) + \chi_i
\end{aligned} \tag{6'}$$

where M_i is an indicator for whether the parental PGI of participant i is missing. The approach requires two assumptions: that there is no assortative mating and that the direct genetic effects and the indirect genetic effects are the same for the participants for whom the parental PGI data are missing and for the participants for whom they are not, both before and after the ROSLA.

We show in Appendix Table I3 that our point estimates are reasonably robust to alternative approaches, namely (i) restricting the sample to those for whom parental PGIs are available and estimating equation (6) and (ii) using the within-siblings variation in the PGI as an instrument for

the PGI. The latter approach has the advantage that it does not rely on the assumption of no assortative mating.

Results

The (reduced-form) results from this analysis are presented in Figure 9. The black circles display the (unconditional) reduced-form estimates of δ_3 in equation (4'), which is the change in the association between the EA PGI and SES caused by the ROSLA. The red triangles and blue diamonds plot the (conditional) reduced-form estimates of θ_3 and θ_5 in equation (6). The red triangles show how the ROSLA changed the direct genetic effects of the PGI on SES. The blue diamonds show how the ROSLA changed the association between the parental PGI and SES. We report results for three SES outcomes: log wages, an indicator for having an annual household income of £31,000 or higher, and an indicator for an annual household income of £52,000 or higher. The left panel shows estimates using a 4-year bandwidth with linear trends. The right panel shows estimates using a 10-year bandwidth and quadratic trends.

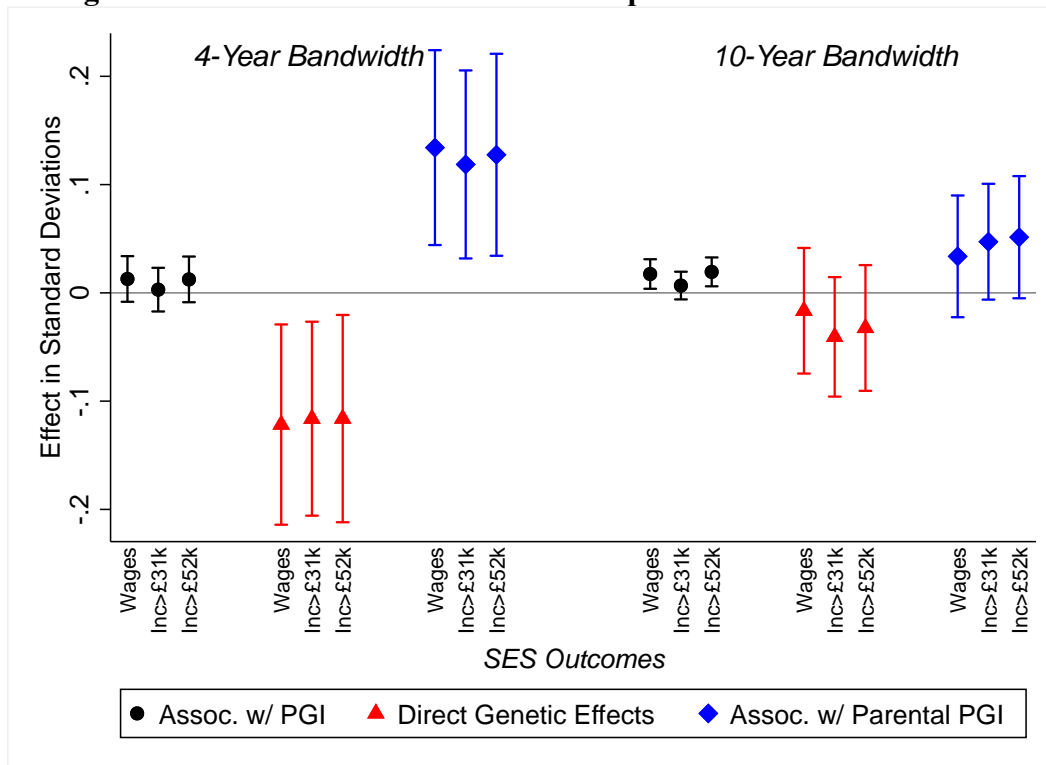
Even with the strategy to account for missingness described above, the analyses in this section are not as well-powered as the other analyses in the paper. This is for two primary reasons. First, the main source of variation that identifies θ_3 and θ_5 (the red triangles and blue diamonds), is the relationship between the children's PGI and the parental PGI, which is only available in about 10% of our sample. Second, the PGI and the parental PGI are highly correlated with a coefficient of correlation of 0.84. As a consequence, the confidence intervals are wide and include some estimates which may be implausibly large.

The black circles in Figure 9 reproduce the results that the ROSLA did not reduce the power of the PGI to predict wages and income. If anything, the policy seemed to have increased it. The red triangles suggest this was not driven by an increase in the direct genetic effects. All point estimates are actually negative, implying that the ROSLA reduced the direct genetic effects. This illustrates well how direct genetic effects are in fact a function of the environment: environmental factors like a change in compulsory schooling laws can change the causal effect of genetic variants. The reduction in the direct genetic effects are statistically significant for the 4-year bandwidth but not for the 10-year.

It seems the ROSLA increased the predictive power of the EA PGI because it made the environmental conditions under which a child grew *more* important for her chances of achieving

socioeconomic success in adulthood. The blue diamonds show that the policy strengthened the association between the parental PGI and (the children's) wages and income.

Figure 9: Effect of ROSLA on Relationship between Genetics and SES



Notes: The figure shows the effect of the ROSLA on the associations with the EA PGI (black circles), on direct genetic effects (red triangles), and on the associations with the parental EA PGI (blue diamonds). The brackets show 95% confidence intervals. The left panel shows estimates for the 4-year bandwidth while the right panel shows the 10-year bandwidth estimates. The EA PGI, the parental EA PGI, and the SES outcomes were all normalized. Wages are imputed based on one's occupation. $N = 46,308$ for 4-year bandwidth and $114,025$ for the 10-year bandwidth.

In Table 3, we exploit the ROSLA to estimate how the returns to schooling vary with the EA PGI. The first four columns show 2SLS results for the 4-year bandwidth. The last four columns show 2SLS results for the 10-year bandwidth. Columns (1), (2), (5), and (6) do not interact the EA PGI with school-leaving age. The even columns control for the parental PGI while the odd columns do not. The corresponding first-stage and reduced-form estimates are shown in Appendix Table I1 and in Appendix Table I2.

Columns (1) and (5) show that the EA PGI predicts wages even when we control for school-leaving age. As expected, this predictive power falls considerably when the parental PGI is added in columns (2) and (6). Nevertheless, the 10-year bandwidth estimate indicates that a one SD

increase in the EA PGI causes wages to increase by 1.5%. In other words, a SD increase in the EA PGI is comparable to increasing education by approximately 0.3 years of schooling.

Table 3: Differences in Returns to Schooling by Genetics and Environment

	Log Wages				Log Wages			
	4-Year Bandwidth				10-Year Bandwidth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PGI * SLA			0.022	-0.137			0.023	-0.011
			[0.012]	[0.060]			[0.007]	[0.026]
Parental PGI * SLA				0.155				0.032
				[0.060]				[0.026]
PGI	0.023	0.004	-0.008	0.189	0.024	0.015	-0.008	0.031
	[0.004]	[0.009]	[0.018]	[0.081]	[0.004]	[0.005]	[0.011]	[0.035]
Parental PGI		0.019		-0.187		0.010		-0.032
		[0.008]		[0.082]		[0.005]		[0.035]
SLA	0.052	0.051	0.059	0.040	0.049	0.049	0.055	0.038
	[0.023]	[0.023]	[0.024]	[0.044]	[0.020]	[0.020]	[0.020]	[0.027]
Observations	46,308	46,308	46,308	46,308	114,025	114,025	114,025	114,025

Notes: The table investigates how the returns to schooling vary with the EA PGI. All estimates are from 2SLS models. The 4-year bandwidth estimates include linear trends. The 10-year bandwidth include quadratic trends. Wages are imputed based on one's occupation. Robust standard errors.

Columns (3) and (7) reproduce results from (the bottom panel of) Table 2, indicating that the EA PGI is associated with higher returns to schooling. Once we condition on the parental PGI in columns (4) and (8), however, the coefficient on the interaction of schooling-leaving age with the PGI turns *negative*, implying that an exogenous increase in the PGI *reduces* the returns to schooling. In contrast, the coefficient on the interaction of schooling-leaving age with the parental PGI is positive, indicating that having parents with higher PGIs (holding the child's genotypes constant) is associated with higher returns to schooling. The 2SLS estimates using a 4-year bandwidth are implausibly large, but they also have wide confidence intervals that include lower returns (in absolute terms). The linear specification with a 4-year bandwidth (column 4) implies that a one SD increase in the EA PGI reduces the returns to schooling by at least 1.9 percentage points.

These results help us understand why advantaged children have higher returns to schooling. Once we condition on the parental PGI, the individual's PGI captures only factors that are causally influenced by that individual's genes. Because we estimate that a higher PGI (conditional on the

parental PGI) reduces the return to schooling, this suggests that genetically-influenced individual characteristics—including innate abilities, skills, or talents—cannot explain why we observe that advantaged children have higher returns. This reduction is consistent with innate skills and talents being substitutes to education in the production of SES (Hause 1972; Welland 1978).

If the positive gradient in returns to schooling cannot be explained by genetics, then it must be due to environmental advantages. This is consistent with our estimate that the returns are higher for children *whose parents* have higher EA PGIs. On average, parents with higher EA PGIs are more educated and earn higher wages. Heckman and co-authors have argued that parents play a crucial role in fostering non-cognitive and cognitive skills of their children (Heckman 2000; Cunha et al. 2006), and that these skills increase the returns to schooling of their offspring (Heckman, Stixrud, and Urzua 2006; Urzua 2006). We discuss in the concluding section some alternative hypotheses for how environmental advantages may increase the returns to schooling.

5. Can Unobservables Explain the Gradient in Returns to Schooling?

Tables 2 and 3 estimate the returns to schooling for compliers, i.e., individuals who, as a consequence of the ROSLA, stayed longer in school than they would have otherwise. Note that compliers in the top vs. bottom tercile of a given proxy may also differ in unobservables. For example, compliers with high EA PGIs might face higher unobservable barriers to education than compliers with low EA PGIs. One question therefore is whether this difference in unobservables may contribute to, or even alternatively explain, the variation in returns to schooling by early-life advantages documented above.

In this section, we investigate this hypothesis using a marginal treatment effects (MTE) framework (Carneiro et al. 2011; Brinch et al. 2017; Kowalski *forthcoming*). This framework requires a binary treatment variable. We consider staying in school until age 16 as our treatment variable – the margin targeted and most affected by the policy – and drop individuals who reported leaving school at age 17 or age 18 (or at older ages). Although this may bias our MTE estimates, only a small fraction of individuals stayed in school until age 17 or age 18 as a result of the ROSLA (see Figure 5) so we anticipate that the bias would be small.

We illustrate the approach by studying the returns to staying in school until age 16 for those in the top and bottom terciles of the distribution of the EA PGI. Results for family size and neighborhood SES can be found in Appendix J. Let Y_1 be the potential log wage if the individual

were to stay in school until age 16, and Y_0 her potential log wage if she were to drop out before age 16. Define the potential outcomes as:

$$Y_1 = \mu_1(\mathbf{X}) + E_1 \text{ and } Y_0 = \mu_0(\mathbf{X}) + E_0 \quad (7)$$

where $\mu_1(\mathbf{x}) = E[Y_1 | \mathbf{X} = \mathbf{x}]$ and $\mu_0(\mathbf{x}) = E[Y_0 | \mathbf{X} = \mathbf{x}]$.

The vector \mathbf{X} includes indicators for the terciles of the EA PGI distribution and birth cohort trends. The return to schooling for the individual is:

$$\beta \equiv Y_1 - Y_0 = \mu_1(\mathbf{x}) - \mu_0(\mathbf{x}) + E_1 - E_0. \quad (8)$$

Let Δ be the net benefit to the individual of staying in school until age 16:

$$\Delta = \Pi(\mathbf{X}, Post) - V \quad (9)$$

where $\Pi(\mathbf{X}, Post)$ is the part of the net benefit that is explained by the observables \mathbf{X} and $Post$ (i.e., the indicator for being born after September 1, 1957) and V is the unobserved (net) cost of staying in school until age 16. The individual decides to stay if $\Delta \geq 0$. An individual for whom $\Pi(\mathbf{x}, Post = 0) < V \leq \Pi(\mathbf{x}, Post = 1)$ is a complier who stayed in school until age 16 as a consequence of the ROSLA when she would not have otherwise.

Define $U \equiv F_V(V)$, where $F_V(V)$ is the cumulative distribution function of V . U can be thought of as the quantile of the individual's unobserved (net) cost of staying in school until age 16. For simplicity, hereafter we refer to U as the “cost” rather than the “quantile of the unobserved net cost” (despite V being more precisely the “cost”). The marginal treatment effect (MTE) is defined by:

$$MTE(\mathbf{x}, u) \equiv E(\beta | \mathbf{X} = \mathbf{x}, U = u) \quad (10)$$

which is the mean return to schooling for individuals with observables \mathbf{x} and an unobserved (net) cost of staying in school until age 16 of u . The goal of this analysis is to compare the MTE of individuals in the bottom and top terciles with the same unobserved costs u .

The local average treatment effect (LATE) is the average MTE among compliers:

$$LATE(\mathbf{x}) = \frac{1}{q(\mathbf{x}) - p(\mathbf{x})} \int_{p(\mathbf{x})}^{q(\mathbf{x})} MTE(\mathbf{x}, u) du \quad (11)$$

where $p(\mathbf{x}) = F_V(\Pi(\mathbf{x}, Post = 0))$ is the fraction of individuals who stayed in school until age 16 before the reform (i.e., the “always-takers”) and $q(\mathbf{x}) = F_V(\Pi(\mathbf{x}, Post = 1))$ is the fraction who stayed until age 16 after the reform (i.e., both the “always-takers” and the “compliers”).

Equation (11) highlights the issue with the LATE estimates reported in Table 2 and Table 3. There are two reasons why the LATE estimates for the bottom and top terciles may be different: (1) the $MTE(\mathbf{x}, u)$ can be different between the two groups, or (2) $p(\mathbf{x})$ and/or $q(\mathbf{x})$ could be different between the two groups. Differences in MTE imply that comparable individuals have different returns to schooling; differences in $p(\mathbf{x})$ or $q(\mathbf{x})$ imply that the LATE estimates of the two groups are based on non-comparable sets of compliers with different (net) unobserved costs. In our context, the fraction of always-takers (i.e., p) is higher in the top PGI tercile than in the bottom tercile (in contrast, the fraction of never-takers is higher in the bottom tercile than in the top tercile). Therefore, the difference in the LATE estimates of the top and bottom terciles may potentially be driven by these differences in unobserved costs between the top-tercile and bottom-tercile compliers.

To address this issue, we propose calculating an “adjusted LATE” which holds constant the unobserved (net) cost of staying in school until age 16:

$$Adjusted\ LATE(\mathbf{x}, p^*, q^*) = \frac{1}{q^* - p^*} \int_{p^*}^{q^*} MTE(\mathbf{x}, u) du \quad (12)$$

where

$$p^* = \max\{p(bottom), p(top)\} = p(top) \text{ and } q^* = \min\{q(bottom), q(top)\} = q(bottom).$$

Individuals with unobserved costs between p^* and q^* would be compliers irrespective of whether they are in the bottom or top tercile of the distribution of the EA PGI.

We can estimate MTEs with only a few assumptions beyond the LATE assumptions (Vytlacil 2002; Imbens and Angrist 1994). These additional assumptions are:

- (1) The cumulative distribution function of V , F_V , must be continuous
- (2) The instrument, $Post$, must be independent of V conditional on \mathbf{x}

(3) $0 < p(\mathbf{x}) < q(\mathbf{x}) < 1$ (i.e., there must be never-takers, compliers, and always-takers)

(4) $E[Y_1 | \mathbf{x}, u]$ and $E[Y_0 | \mathbf{x}, u]$ are linear functions of the (net) unobserved cost u

Note that because V was defined as the portion of the net benefit unexplained by $Post$ and \mathbf{x} , (2) is true by construction. Details on all of our assumptions and the derivation of our MTE estimator can be found in Appendix J.

By Assumption (4), we can express the MTE of staying in school till age 16 as:

$$MTE(\mathbf{x}, u) = \lambda_0(\mathbf{x}) + [\lambda_1(\mathbf{x}) \times u]. \quad (10')$$

In this case, LATE and the adjusted LATE are given by:

$$LATE(\mathbf{x}) = MTE\left(\mathbf{x}, \frac{p(\mathbf{x}) + q(\mathbf{x})}{2}\right) \quad (11')$$

$$Adjusted\ LATE(\mathbf{x}, p^*, q^*) = MTE\left(\mathbf{x}, \frac{p^* + q^*}{2}\right) \quad (12')$$

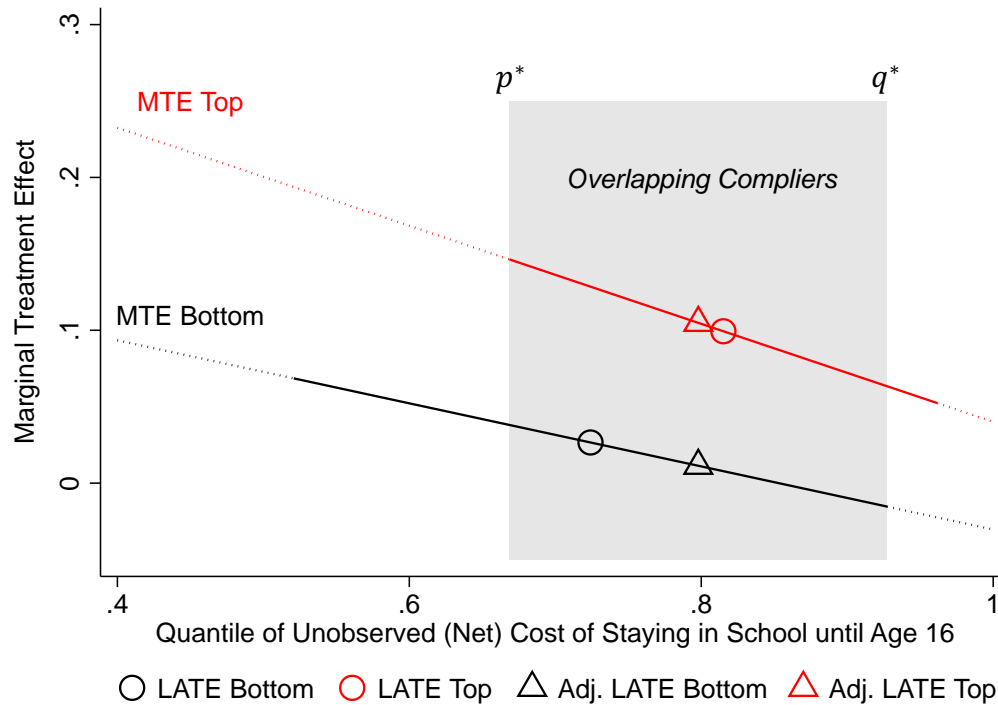
Appendix J shows how one can identify $\lambda_0(\mathbf{x})$ and $\lambda_1(\mathbf{x})$ for the top and bottom terciles of the PGI distribution.

Figure 10 shows estimates of the MTE, LATE, and the adjusted LATE, separately for the bottom and top terciles. The black line shows the MTE for the bottom tercile. The solid segment of the red line indicates the MTEs for the set of compliers in the bottom tercile. Similarly, the solid segment of the black line indicates the MTEs for the set of compliers in the top tercile. The area shaded in gray indicates the “overlapping compliers” with unobserved cost ranging from p^* to q^* . The circles show the LATE estimates for the bottom (red) and top (black) terciles. The triangles show the adjusted LATEs.⁴⁰

Notice how the MTE lines slope downwards, such that the mean return to schooling is actually decreasing in the unobserved (net) cost. The “marginal complier” (i.e., the complier who is indifferent between staying until 16 or dropping out) in the top tercile has a higher unobserved (net) cost of staying in school until age 16 than the marginal complier in the bottom tercile. As a consequence, the top-bottom difference in the adjusted LATEs is larger than the top-bottom difference in the “naïve LATEs.”

⁴⁰ Due to the linearity assumption, these points are located at the middle of each relevant complier range.

Figure 10: Marginal Treatment Effects of Top and Bottom Terciles of EA PGI Distribution



Notes: The figure shows estimates of the MTE, LATE, and the adjusted LATE, separately for the bottom and top terciles of the EA PGI distribution. The black line shows the MTE for the bottom tercile. The solid segment of the black line flags the compliers in the bottom tercile. The red line shows the MTE for the top tercile. Its solid segment flags the compliers in the top tercile. The shaded area flags “overlapping compliers” with unobserved costs ranging from p^* to q^* . The black circle and the black triangle show respectively the LATE and adjusted LATE for the bottom tercile. The red circle and the red triangle shows respectively the LATE and adjusted LATE the top tercile. $N = 70,670$ using a 10-year bandwidth.

Appendix Table J1 shows that the LATE estimates of the returns of staying in school until age 16 is 7.2 percentage points larger for the top tercile than for the bottom tercile. In comparison, the adjusted LATE estimates of these returns is 9.4 percentage points larger for the top tercile. We can reject at the 1% level that the difference in adjusted LATEs is zero. Appendix Table J1 shows that the same conclusion is reached if we use a 4-year bandwidth or look at the other two proxies for early-life advantage, family size or neighborhood SES. We therefore conclude that the gradient in returns to schooling documented above cannot be explained by differences in unobservable characteristics between advantaged and disadvantaged compliers. If anything, the gradient *increases* once we control for such differences.

One of the model’s assumption is that there are at least some never-takers (assumption 3 above). About 5% of those born after September 1, 1957 report having left school before age 16,

which means that they would not have abided by the ROSLA.⁴² Although the MTE would not be identified if there were 100% compliance with the policy, we show in Appendix J that one would still be able to calculate how large the slope of the bottom tercile's MTE would have to be in order for the top-bottom difference in LATE to be solely driven by the difference in the unobserved (net) costs of compliers in the bottom and top terciles. These calculations imply, for example, that the return of staying in school until age 16 of the average complier would have to be 50 percentage points *larger* than the return of the average always-taker. We find it highly implausible that compliers would have dropped out before age 16 in the absence of the ROSLA if their returns to staying in school were so much larger.

6. Robustness

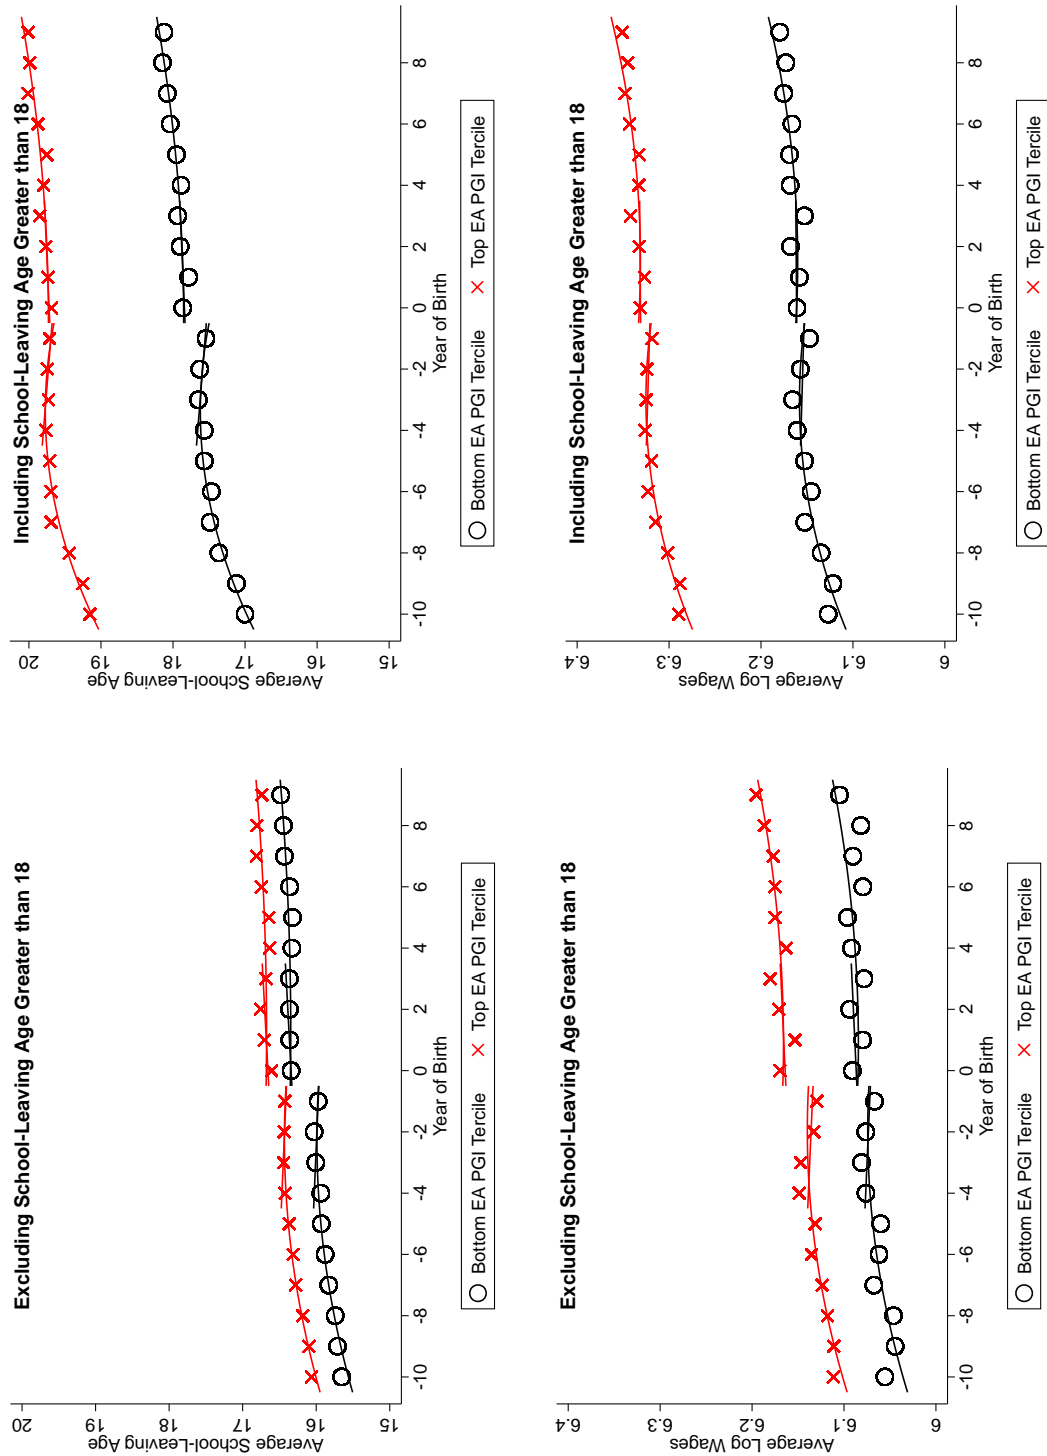
In this section, we report on a series of exercises to investigate the robustness of our results. We show how the results (do not) change when we (a) Include in the sample those who stayed in school past age 18; (b) Allow different terciles to have different birth cohort trends; (c) Vary the regression discontinuity bandwidth; (d) Include controls; and (e) Use income or the Townsend Deprivation Index as the SES outcome.

A. Exclusion of Those who Stayed in School Past Age 18

Figure 11 assesses the sensitivity to excluding individuals who stayed in school past age 18. The left column shows results for our main sample, which excludes these individuals. The right column shows results for the full sample, which includes those who stayed in school past age 18. The top row shows average school-leaving age by year of birth, separately for the bottom and top terciles of the EA PGI distribution. The bottom row shows average log wages by year of birth.

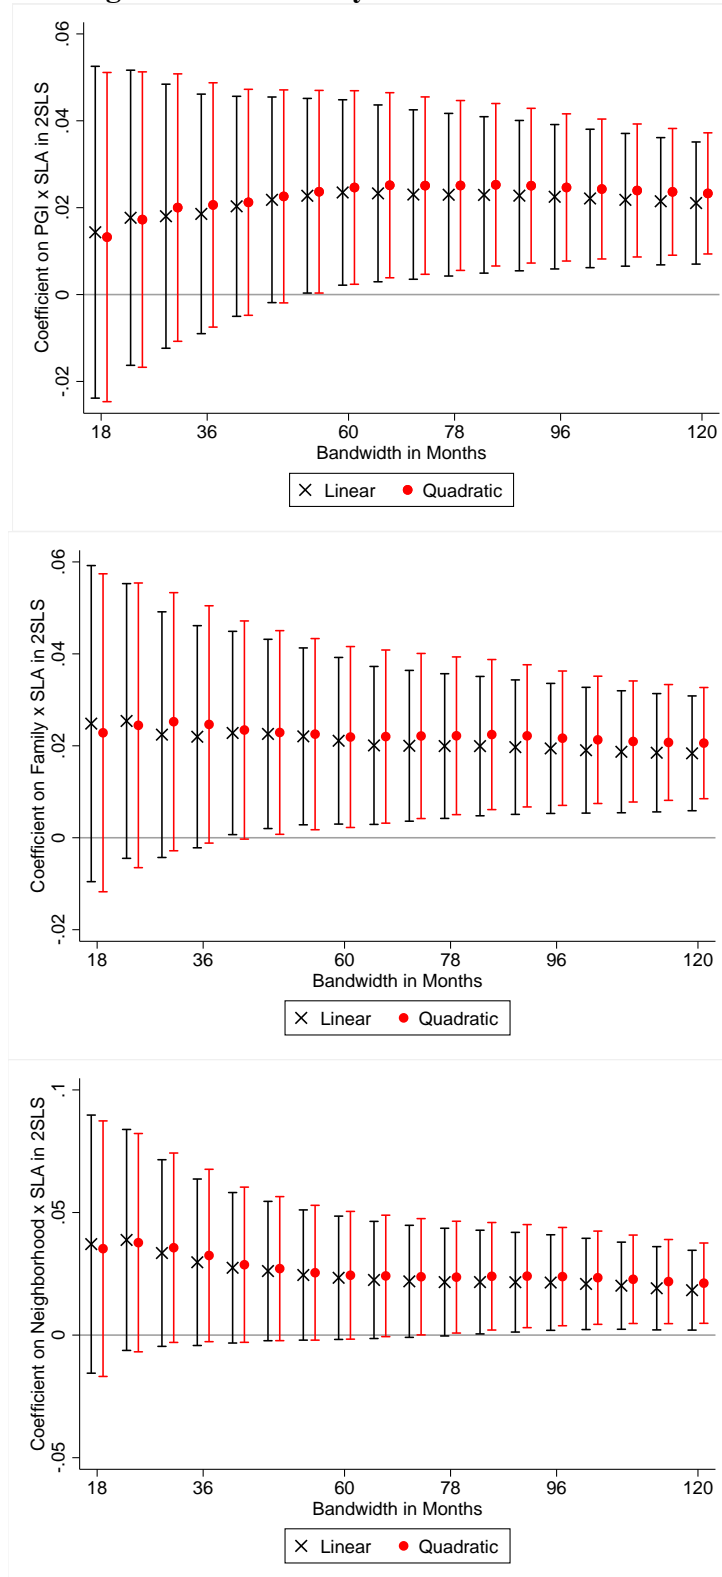
⁴² For students born in summer months, school may have ended when they were still 15 but they would have been 16 by the time the next academic year started. Such a student may have reported a school-leaving age of 15 but would have been in compliance with the 1972 ROSLA. About 2/3 of the potential never-takers were born in summer months.

Figure 11: Including UKB Participants who Stayed in School Past Age 18



Notes: The figure assesses the sensitivity to excluding individuals who stayed in school past age 18. The left column shows results for our main sample, which excludes these individuals. The right column shows results for the full sample, which includes those who stayed in school past age 18. The top row shows average school-leaving age by year of birth, separately for the bottom and top tertiles of the EA PGI distribution. The bottom row shows average log wages by year of birth. $N = 42,731$ (excluding SLA > 18 with 4-year bandwidth); $105,693$ (excluding SLA > 18 with 10-year bandwidth); $79,632$ (including SLA > 18 with 4-year bandwidth); and $196,004$ (including SLA > 18 with 10-year bandwidth). Wages are imputed based on one's occupation.

Figure 12: Sensitivity to Bandwidth Choice



Notes: The figure investigates how the 2SLS estimates (from continuous specification) vary with the bandwidth size (in months). It shows the coefficient on the interaction of school-leaving age with the EA PGI (top panel), family size (middle panel), and neighborhood SES (bottom panel). Black Xs show estimates using linear trends. Red circles show estimates using quadratic trends. Wages are imputed based on one's occupation. N varies from 17,783 (18-month bandwidth) to 114,025 (120-month bandwidth).

Figure 11 shows that the main conclusions are invariant to the sample restriction. Both the left and right columns show that, while the ROSLA reduced the difference in schooling between the top and bottom terciles, there is no evidence of a reduction in the wage gap between them. The regression discontinuities are obviously smaller in the right column because these estimates include a lot of participants who were not affected by the ROSLA. That is particularly true for the top tercile because the fraction of participants staying in school past age 18 was higher in the top tercile than in the bottom tercile. Similar figures are shown in Appendix N for family size and neighborhood SES.

Table 4: Robustness

	First Stage			2SLS				
				<i>School-Leaving Age</i>			<i>Log Wages</i>	
						<i>≥ £31,000</i>	<i>Index</i>	
PGI * Post	-0.042	-0.042	-0.044					
	[0.007]	[0.007]	[0.007]					
Family * Post	-0.080	-0.080	-0.077					
	[0.007]	[0.007]	[0.007]					
Neighborhood * Post	-0.043	-0.043	-0.035					
	[0.007]	[0.007]	[0.007]					
Post	0.331	0.290	0.288					
	[0.018]	[0.018]	[0.019]					
PGI * SLA				0.022	0.020	0.019	0.023	0.052
				[0.007]	[0.006]	[0.007]	[0.010]	[0.019]
Family * SLA				0.019	0.021	0.020	0.037	0.060
				[0.006]	[0.006]	[0.006]	[0.009]	[0.017]
Neighborhood * SLA				0.019	0.019	0.016	-0.013	0.010
				[0.008]	[0.008]	[0.008]	[0.010]	[0.021]
SLA				0.061	0.065	0.070	0.137	0.091
				[0.022]	[0.022]	[0.024]	[0.034]	[0.066]
Controls?		✓	✓		✓	✓	✓	✓
PCs?			✓			✓	✓	✓
Observations	114,025	114,025	114,025	114,025	114,025	114,025	133,086	155,489

Notes: The first six columns of the table assess the robustness to including controls. All columns except for columns 1 and 4 include the following controls: male, age, age squared, male × age, male × age squared, and indicator variables for the calendar month of birth. To address concerns about population stratification, columns 3 and 6 include in addition the first twenty principal components (PCs) of the participant’s genetic matrix and the interactions of these PCs with either the indicator for being born after September 1, 1957 or with school-leaving age. The dependent variable in columns 1-6 is log wages, which are imputed based on one’s occupation. To investigate concerns about missing wages, columns 7 and 8 use as the dependent variable an indicator for having an annual household income of £31,000 or more and the Townsend Deprivation Index (reverse coded). Robust standard errors between brackets.

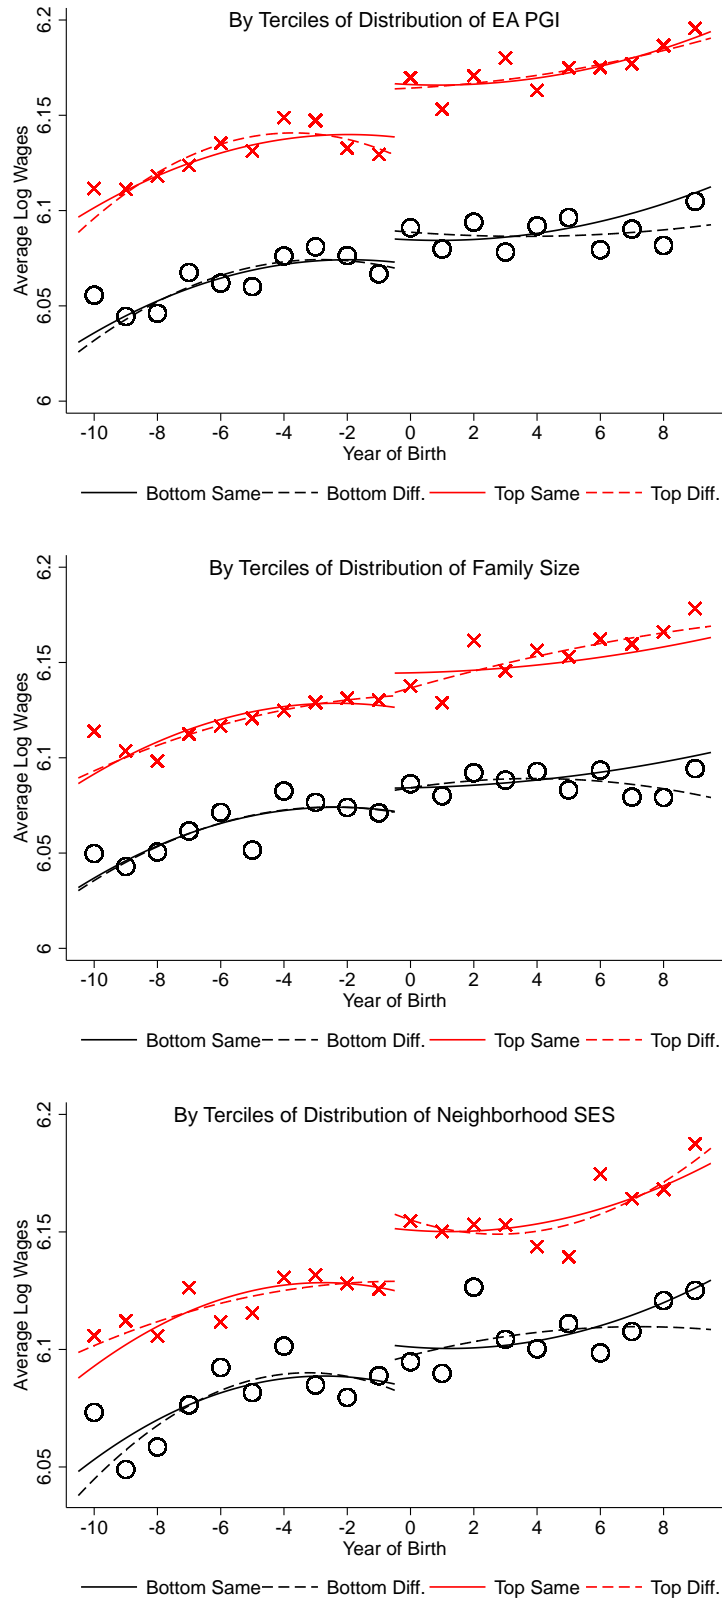
B. Bandwidth Choice

Figure 12 assesses the sensitivity of the results to the choice of bandwidth. It shows two-stage least square estimates for different bandwidths (measured in months). The markers display the coefficient on the interaction of one of the proxies with school-leaving age (i.e., η_3 in the continuous specification of equation (5)) with 95% confidence intervals around it. The black Xs are estimated using linear trends. The red circles are estimated using quadratic trends. The figure shows that the point estimates are remarkably similar across bandwidths and trend specifications. They are not statistically different from zero at the 5% confidence level for narrower bandwidths, but this is expected because the confidence intervals increase as the bandwidth shrinks. The corresponding figures for the first-stage estimates and for the reduced-form estimates are shown in Appendix G.

C. Allowing Birth Cohort Trends to Vary with Proxies of Early-Life Advantage

To maximize statistical power, the previous analyses assumed that the birth cohort trends were invariant to early-life advantages. Figure 13 assesses the sensitivity to this assumption. It shows two sets of quadratic trends for the 10-year bandwidth. The solid curves show the birth cohort trends for log wages under this assumption. The dashed curves show comparable trends when the top and bottom terciles of a given proxy are allowed to have distinct birth cohort trends (the middle tercile also has its own birth cohort trend; it is not shown in the figure for ease of exposition). The two sets of trends are in general quite similar. They are especially similar for those born around 1957, which is what relevant for the regression discontinuity. Appendix Figure M1 and Appendix Figure M2 show birth cohort trends for school-leaving age and for qualifications, yielding the same conclusion. Appendix Figure M4, M5, and M6 show that the results also hold for the 4-year bandwidth with linear trends. Appendix Table M1, M2, M3, and M4 show corresponding estimates. Overall, the point estimates of the top-bottom difference in returns to schooling seem to get even larger – with the exception of family size (see Appendix Table M4). Nevertheless, they lose statistical significance because standard errors typically double or triple. Importantly, the p-values reported in the tables show that we cannot reject at the 5% level that the different terciles have the same birth cohort trends.

Figure 13: Different Birth Cohort Trends for Bottom and Top Terciles



Notes: The figure assesses the sensitivity of the results to the assumption that the birth cohort trends were invariant to early-life advantages. The solid curves show 10-year birth cohort trends for log wages under this assumption. The dashed curves show comparable trends when the top and

bottom terciles are allowed to have distinct birth cohort trends (the middle tercile also has its own birth cohort trend; it is not shown for ease of exposition). $N = 105,693$ (EA PGI); 112,395 (family size); and 109,177 (neighborhood SES). Wages are imputed based on one's occupation.

D. Controls

Table 4 investigates whether the results are robust to the inclusion of controls. For conciseness, we only report results for the continuous specification using a 10-year bandwidth with quadratic trends. The first column of Table 4 reproduces the first-stage benchmark (which corresponds to the last column, bottom panel of Table 1) while the fourth column reproduces the 2SLS benchmark (which corresponds to the last column, bottom panel of Table 2). The second and fourth columns of Table 4 adds the following controls: male, age, age squared, male \times age, male \times age squared, and indicator variables for the calendar month of birth. To address concerns about population stratification (see Section 3.A), the third and sixth columns include, besides the controls previously mentioned, the first twenty principal components (PCs) of the participant's genetic matrix and the interactions of these PCs with either the indicator for being born after September 1, 1957 (in the first-stage) or with school-leaving age (in the 2SLS). The first-stage results in columns 2 and 3 are not sensitive to these inclusions. Nor are the 2SLS results shown in columns 5 and 6 of Table 4. That said, there is evidence that some social stratification in a PGI may remain between historically poorer and richer regions of the UK even after controlling for principal components (Abdellaoui et al. 2019).

E. Alternative SES Measures

Finally, we investigate the concern that data on wages are available only for those UKB participants who reported an occupation. The last two columns of Table 4 report two-stage least square estimates that use as dependent variables either an indicator for having an annual household income of £31,000 or more or a Townsend Deprivation Index (reverse coded). While wages are missing for 26.8% of our sample of 155,806 UKB participants, income is missing for 14.5% and the Townsend Index for 0.1%. The Townsend Deprivation Index is a measure of the material deprivation of the neighborhoods in which study participants lived.⁴³ Overall, we reach similar

⁴³ The Townsend Deprivation Index is constructed from four rates measured at the neighborhood level: (1) unemployment, (2) non-car ownership, (3) non-home ownership, and (4) household overcrowding. Rates were estimated for each output area using the 2001 Census. Participants were assigned the scores of the output areas where their residential postcodes were located.

conclusions irrespective of whether we use occupational wages, household income (see also Appendix H), or the Townsend index. The exception is that, when we use these alternative outcomes, the returns to schooling no longer vary with neighborhood SES. These results are consistent with the evidence shown in the Appendix E that the fraction missing wages and the fraction missing household income are smooth around the September 1, 1957 birthday cutoff.

7. Conclusion

Many consider education to be the “great equalizer” that “levels the playing field,” giving children from different backgrounds similar opportunities of economic success. We investigated whether education weakens the relationship between early-life advantages and later-life SES. In particular, we examined whether the ROSLA, a 1972 compulsory schooling reform in the UK that increased secondary education, reduced disparities between disadvantaged and advantaged children. Besides early, favorable family and neighborhood circumstances, we argue that the genes a child inherits also represent a source of individual-level advantages. We find that the ROSLA was successful in reducing the differences in education between children from advantaged and disadvantaged backgrounds. Meghir and Palme (2005) and Aakvik et al. (2010) find similar results for reforms in Sweden and Norway.

Our results indicate, however, that, when it comes to later-life SES, the UK reform was not a “great equalizer.” While the ROSLA increased wages for individuals from all backgrounds, it did not reduce the wage gaps between them. This is consistent with the idea that education might be more effective in reducing poverty than in promoting mobility (Cameron and Heckman 2001). Our finding contrasts with the impacts of the Swedish reform, a similar reform in Finland, and a more recent reform in Norway. The Swedish reform increased the wages of children of less-educated fathers and decreased the wages of children of more-educated fathers (Meghir and Palme 2005); the Finnish reform reduced the correlation between the earnings of fathers and sons (Pekkarinen et al. 2009); and the Norwegian reform reduced the gap in earnings between disadvantaged and advantaged children (Bertrand, Mogstad and Mountjoy 2020).

Different from the reforms in Sweden, Norway and Finland, the key elements of the British school system did not change with the 1972 ROSLA. This distinction enables us to estimate how returns to schooling vary with children’s background and to identify the role that such differences may play in the intergenerational transmission of SES. We find that advantaged children had

higher returns to schooling which explains why, despite reducing educational disparities, the UK reform did not weaken the relationship between early advantages and wages. In contrast, the Swedish, Finnish, and Norwegian reforms shifted a large fraction of disadvantaged children from a vocational stream with lower returns to schooling to a more academic stream with higher returns.⁴⁴ The contrast between these two types of reforms suggests that differences in the returns to schooling may be an important challenge to increasing intergenerational mobility (Solon 2004).

We then turn to the question of why advantaged children saw higher returns following the UK reform. For example, did the compulsory schooling unlock the individual, untapped potential of children who would have otherwise dropped out at age 15? Did it make it easier for children from advantaged families to leverage their resources and connections to get further ahead? We exploit family-based random genetic variation to tackle this question. Our results reject the hypothesis that individual-level, genetically-influenced characteristics (including abilities, skills, or talents) explain why advantaged children had higher returns, suggesting instead that it reflects environmental channels that might represent unequal opportunities.

School quality is an example of such a potential channel. We found higher returns to schooling for children born in more-educated neighborhoods, which may have, on average, higher-quality schools than less-educated neighborhoods. In the US, returns to schooling are larger at regions with higher-quality schools (Card and Krueger 1992; DeCicca and Krashinsky 2020). On the other hand, the limited existing evidence for the UK suggests that the returns to schooling in England and Wales do not vary with school quality (Campbell 2001). An alternative hypothesis is that more educated neighborhoods have more economic activity. Due to frictions in migration, children born in less educated neighborhoods may not benefit as much from additional education since there are fewer high-paying jobs available where they live (Bergman et al. 2019). Finally, such environmental advantages could play out at the family-level. Heckman and co-authors have argued that parents play a crucial role in fostering non-cognitive and cognitive skills of their children (Heckman 2000; Cunha et al. 2006), and that these skills increase the returns to schooling of their offspring (Heckman, Stixrud, and Urzua 2006; Urzua 2006).

⁴⁴ Starting in seventh grade, students in Sweden who attended the old school system were tracked into an academic or a vocational stream based on their grades. The reform abolished this selection, giving students the option to choose between three different streams. In Finland, students could also choose between an academic and a vocational track. The reform postponed this choice from age 11 to age 16. In Norway, the reform aimed to improve the quality of the vocational track by integrating more general education into the curriculum and offering vocational students a pathway to college.

More research is needed to understand which family- and neighborhood-advantages matter in this case and, more generally, whether our results hold in other contexts and for other margins of education. While compulsory schooling alone does not appear to decrease wage disparities and promote intergenerational mobility, other types of policies instead of or in conjunction with compulsory schooling may be more effective. Our results imply that policies designed to increase intergenerational mobility should target environmental barriers that reduce the returns to schooling of disadvantaged children.

References

- Aakvik, Arild, Kjell G. Salvanes, and Kjell Vaage. "Measuring heterogeneity in the returns to education using an education reform." *European Economic Review* 54, no. 4 (2010): 483-500.
- Abdellaoui, Abdel, David Hugh-Jones, Loic Yengo, Kathryn E. Kemper, Michel G. Nivard, Laura Veul, Yan Holtz et al. "Genetic correlates of social stratification in Great Britain." *Nature human behaviour* 3, no. 12 (2019): 1332-1342.
- Allen, Naomi, Cathie Sudlow, Paul Downey, Tim Peakman, John Danesh, Paul Elliott, John Gallacher et al. "UK Biobank: Current status and what it means for epidemiology." *Health Policy and Technology* 1, no. 3 (2012): 123-126.
- Almond, Douglas, Janet Currie, and Valentina Duque. "Childhood circumstances and adult outcomes: Act II." *Journal of Economic Literature* 56, no. 4 (2018): 1360-1446.
- Altonji, Joseph G., and Thomas A. Dunn. "Using siblings to estimate the effect of school quality on wages." *The Review of Economics and Statistics* (1996): 665-671.
- Andreoli, Francesco, Markus Jantti, and Arnaud Lefranc. "Does Comprehensive Compulsory Schooling Equalize Opportunities? Evidence from the Swedish Compulsory Schooling Reform," 2020.
- Angrist, Joshua, Victor Lavy, and Analia Schlosser. "Multiple experiments for the causal link between the quantity and quality of children." *Journal of Labor Economics* 28, no. 4 (2010): 773-824.
- Ashenfelter, Orley, and Cecilia Rouse. "Income, schooling, and ability: Evidence from a new sample of identical twins." *The Quarterly Journal of Economics* 113, no. 1 (1998): 253-284.
- Banks, James, and Fabrizio Mazzonna. "The effect of education on old age cognitive abilities: evidence from a regression discontinuity design." *The Economic Journal* 122, no. 560 (2012): 418-448.

- Banks, James, Leandro S. Carvalho, and Francisco Perez-Arce. "Education, decision making, and economic rationality." *Review of Economics and Statistics* 101, no. 3 (2019): 428-441.
- Barcellos, Silvia H., Leandro S. Carvalho, and Patrick Turley. "Education can reduce health differences related to genetic risk of obesity." *Proceedings of the National Academy of Sciences* 115, no. 42 (2018): E9765-E9772.
- Barcellos, Silvia H., Leandro S. Carvalho, and Patrick Turley. "Distributional effects of education on health." Forthcoming at the *Journal of Human Resources*.
- Barth, Daniel, Nicholas W. Papageorge, and Kevin Thom. "Genetic endowments and wealth inequality." *Journal of Political Economy* 128, no. 4 (2020): 1474-1522.
- Belsky, Daniel W., Benjamin W. Domingue, Robbee Wedow, Louise Arseneault, Jason D. Boardman, Avshalom Caspi, Dalton Conley et al. "Genetic analysis of social-class mobility in five longitudinal studies." *Proceedings of the National Academy of Sciences* 115, no. 31 (2018): E7275-E7284.
- Benabou, Roland. "Unequal societies: Income distribution and the social contract." *American Economic Review* 90, no. 1 (2000): 96-129.
- Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer. *Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice*. No. w26164. National Bureau of Economic Research, 2019.
- Bertrand, Marianne, Magne Mogstad, and Jack Mountjoy. "Improving Educational Pathways to Social Mobility: Evidence from Norway's "Reform 94"." Forthcoming at *Journal of Labor Economics*.
- Biroli, Pietro, and Christian Lukas Zünd. "Genes, Pubs, and Drinks: Gene-environment interplay and alcohol licensing policy in the UK." (2020).
- Biroli, Pietro, and Laura Zwysig. "Moral Hazard Heterogeneity: Genes and Health Insurance Influence Smoking after a Health Shock." (2021).
- Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. "The more the merrier? The effect of family size and birth order on children's education." *The Quarterly Journal of Economics* 120, no. 2 (2005): 669-700.
- Black, Sandra E., Paul J. Devereux, Petter Lundborg, and Kaveh Majlesi. "Poor little rich kids? The role of nature versus nurture in wealth and other economic outcomes and behaviours." *The Review of Economic Studies* 87, no. 4 (2020): 1683-1725.
- Bowles, Samuel, and Herbert Gintis. "The inheritance of inequality." *Journal of Economic Perspectives* 16, no. 3 (2002): 3-30.
- Brinch, Christian N., Magne Mogstad, and Matthew Wiswall. "Beyond LATE with a discrete instrument." *Journal of Political Economy* 125, no. 4 (2017): 985-1039.
- Brumpton, Ben, Eleanor Sanderson, Karl Heilbron, Fernando Pires Hartwig, Sean Harrison, Gunnhild Åberge Vie, Yoonsu Cho et al. "Avoiding dynastic, assortative mating, and population

stratification biases in Mendelian randomization through within-family analyses." *Nature communications* 11, no. 1 (2020): 1-13.

Cameron, Stephen V., and James J. Heckman. "The dynamics of educational attainment for black, hispanic, and white males." *Journal of political Economy* 109, no. 3 (2001): 455-499.

Campbell, David. *Rates of return to schooling and the quality of education in England and Wales*. No. 01, 15. Department of Economics Discussion Paper, 2001.

Card, David. "Estimating the return to schooling: Progress on some persistent econometric problems." *Econometrica* 69, no. 5 (2001): 1127-1160.

Card, David, and Alan B. Krueger. "Does school quality matter? Returns to education and the characteristics of public schools in the United States." *Journal of political Economy* 100, no. 1 (1992): 1-40.

Carneiro, Pedro, James J. Heckman, and Edward J. Vytlačil. "Estimating marginal returns to education." *American Economic Review* 101, no. 6 (2011): 2754-81.

Case, Anne, and Christina Paxson. "Stature and status: Height, ability, and labor market outcomes." *Journal of political Economy* 116, no. 3 (2008): 499-532.

Caspi, Avshalom, Joseph McClay, Terrie E. Moffitt, Jonathan Mill, Judy Martin, Ian W. Craig, Alan Taylor, and Richie Poulton. "Role of genotype in the cycle of violence in maltreated children." *Science* 297, no. 5582 (2002): 851-854.

Caspi, Avshalom, Karen Sugden, Terrie E. Moffitt, Alan Taylor, Ian W. Craig, HonaLee Harrington, Joseph McClay et al. "Influence of life stress on depression: moderation by a polymorphism in the 5-HTT gene." *Science* 301, no. 5631 (2003): 386-389.

Chan, Tak Wing, Morag Henderson, and Rachel Stuchbury. "Family size and educational attainment in England and Wales." *Population studies* 73, no. 2 (2019): 165-178.

Cheesman, Rosa, Avina Hunjan, Jonathan RI Coleman, Yasmin Ahmadzadeh, Robert Plomin, Tom A. McAdams, Thalia C. Eley, and Gerome Breen. "Comparison of adopted and nonadopted individuals reveals gene–environment interplay for education in the UK Biobank." *Psychological science* 31, no. 5 (2020): 582-591.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. "The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment." *American Economic Review* 106, no. 4 (2016): 855-902.

Chetty, Raj, and Nathaniel Hendren. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics* 133, no. 3 (2018a): 1107-1162.

- Chetty, Raj, and Nathaniel Hendren. "The impacts of neighborhoods on intergenerational mobility II: County-level estimates." *The Quarterly Journal of Economics* 133, no. 3 (2018b): 1163-1228.
- Clark, Damon, and Heather Royer. *The effect of education on adult health and mortality: Evidence from Britain*. No. w16013. National Bureau of Economic Research, 2010.
- Clark, Damon, and Heather Royer. "The effect of education on adult mortality and health: Evidence from Britain." *American Economic Review* 103, no. 6 (2013): 2087-2120.
- Conti, Gabriella, and James J. Heckman. "Understanding the early origins of the education–health gradient: A framework that can also be applied to analyze gene–environment interactions." *Perspectives on Psychological Science* 5, no. 5 (2010): 585-605.
- Crowther, G. "The Crowther report, 15 to 18: A report of the Central Advisory Council for Education (England)." (1959).
- Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy V. Masterov. "Interpreting the evidence on life cycle skill formation." *Handbook of the Economics of Education* 1 (2006): 697-812.
- Currie, Janet, and Douglas Almond. "Human capital development before age five." In *Handbook of labor economics*, vol. 4, pp. 1315-1486. Elsevier, 2011.
- D'onofrio, Brian M., Benjamin B. Lahey, Eric Turkheimer, and Paul Lichtenstein. "Critical need for family-based, quasi-experimental designs in integrating genetic and social science research." *American journal of public health* 103, no. S1 (2013): S46-S55.
- Davies, Neil M., Matt Dickson, George Davey Smith, Gerard J. Van Den Berg, and Frank Windmeijer. "The causal effects of education on health outcomes in the UK Biobank." *Nature human behaviour* 2, no. 2 (2018): 117-125.
- DeCicca, Philip, and Harry Krashinsky. *Do Differences in School Quality Generate Heterogeneity in the Causal Returns to Education?*. No. w27089. National Bureau of Economic Research, 2020.
- Deschênes, Olivier. "Estimating the effects of family background on the return to schooling." *Journal of Business & Economic Statistics* 25, no. 3 (2007): 265-277.
- Dickson, Matt, and Sarah Smith. "What determines the return to education: an extra year or a hurdle cleared?." *Economics of education review* 30, no. 6 (2011): 1167-1176.
- Domingue, Benjamin W., and Jason Fletcher. "Separating measured genetic and environmental effects: Evidence linking parental genotype and adopted child outcomes." *Behavior genetics* 50, no. 5 (2020): 301-309.
- Fagereng, Andreas, Magne Mogstad, and Marte Rønning. "Why do wealthy parents have wealthy children?." *Journal of Political Economy* 129, no. 3 (2021): 703-756.

Fletcher, Jason M., and Dalton Conley. "The challenge of causal inference in gene–environment interaction research: Leveraging research designs from the social sciences." *American journal of public health* 103, no. S1 (2013): S42-S45.

Fletcher, Jason, Yuchang Wu, Zijie Zhao, and Qiongshi Lu. "The production of within-family inequality: Insights and implications of integrating genetic data." *BioRxiv* (2020).

Grenet, Julien. "Is extending compulsory schooling alone enough to raise earnings? Evidence from French and British compulsory schooling laws." *The Scandinavian Journal of Economics* 115, no. 1 (2013): 176-210.

Harden, K. Paige, Benjamin W. Domingue, Daniel W. Belsky, Jason D. Boardman, Robert Crosnoe, Margherita Malanchini, Michel Nivard, Elliot M. Tucker-Drob, and Kathleen Mullan Harris. "Genetic associations with mathematics tracking and persistence in secondary school." *NPJ science of learning* 5, no. 1 (2020): 1-8.

Hause, John C. "Earnings profile: Ability and schooling." *Journal of political economy* 80, no. 3, Part 2 (1972): S108-S138.

Heckman, James J. "Policies to foster human capital." *Research in economics* 54, no. 1 (2000): 3-56.

Heckman, James J., Jora Stixrud, and Sergio Urzua. "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior." *Journal of Labor economics* 24, no. 3 (2006): 411-482.

Heckman, James J., John Eric Humphries, and Gregory Veramendi. "Returns to education: The causal effects of education on earnings, health, and smoking." *Journal of Political Economy* 126, no. S1 (2018): S197-S246.

Howe, Laurence, David Evans, Gibran Hemani, George Davey Smith, and Neil Davies. "Evaluating indirect genetic effects of siblings using singletons." (2021).

Imbens, GW, and JD Angrist. "Identification and estimation of local average treatment effects." *Econometrica* 62, no. 2 (1994): 467-475.

Jencks, Christopher. "Who gets ahead? The determinants of economic success in America." (1979).

Johnson, Rebecca A., Ramina Sotoudeh, and Dalton Conley. "Polygenic scores for plasticity: A new tool for studying gene-environment interplay." *bioRxiv* (2020).

Kong, Augustine, Gudmar Thorleifsson, Michael L. Frigge, Bjarni J. Vilhjalmsson, Alexander I. Young, Thorgeir E. Thorgeirsson, Stefania Benonisdottir et al. "The nature of nurture: Effects of parental genotypes." *Science* 359, no. 6374 (2018): 424-428.

Kong, Augustine, Stefania Benonisdottir, and Alexander I. Young. "Family analysis with Mendelian imputations." *BioRxiv*(2020).

- Kowalski, Amanda E. "Reconciling seemingly contradictory results from the Oregon health insurance experiment and the Massachusetts health reform." Forthcoming in the *Review of Economics and Statistics*.
- Kweon, Hyeokmoon, Casper Burik, Richard Karlsson Linnér, Ronald De Vlaming, Aysu Okbay, Daphne Martschenko, Kathryn Harden Harden, Thomas A. DiPrete, and Philipp Koellinger. "Genetic Fortune: Winning or Losing Education, Income, and Health." (2020).
- Lee, David S., and David Card. "Regression discontinuity inference with specification error." *Journal of Econometrics* 142, no. 2 (2008): 655-674.
- Lee, James J., Robbee Wedow, Aysu Okbay, Edward Kong, Omeed Maghzian, Meghan Zacher, Tuan Anh Nguyen-Viet et al. "Gene discovery and polygenic prediction from a genome-wide association study of educational attainment in 1.1 million individuals." *Nature genetics* 50, no. 8 (2018): 1112-1121.
- Lee, David L., Justin McCrary, Marcelo J. Moreira, and Jack Porter. "Valid t-ratio Inference for IV." *arXiv preprint arXiv:2010.05058* (2020).
- Loury, Glenn C. "Intergenerational transfers and the distribution of earnings." *Econometrica: Journal of the Econometric Society* (1981): 843-867.
- Martin, Alicia R., Masahiro Kanai, Yoichiro Kamatani, Yukinori Okada, Benjamin M. Neale, and Mark J. Daly. "Clinical use of current polygenic risk scores may exacerbate health disparities." *Nature genetics* 51, no. 4 (2019): 584-591.
- Meghir, Costas, and Mårten Palme. "Educational reform, ability, and family background." *American Economic Review* 95, no. 1 (2005): 414-424.
- Mogstad, Magne, and Matthew Wiswall. "Testing the quantity–quality model of fertility: Estimation using unrestricted family size models." *Quantitative Economics* 7, no. 1 (2016): 157-192.
- Morris, Tim T., Neil M. Davies, Gibran Hemani, and George Davey Smith. "Population phenomena inflate genetic associations of complex social traits." *Science Advances* 6, no. 16 (2020): eaay0328.
- Okbay, Aysu, Jonathan P. Beauchamp, Mark Alan Fontana, James J. Lee, Tune H. Pers, Cornelius A. Rietveld, Patrick Turley et al. "Genome-wide association study identifies 74 loci associated with educational attainment." *Nature* 533, no. 7604 (2016): 539-542.
- Oreopoulos, Philip. "Estimating average and local average treatment effects of education when compulsory schooling laws really matter." *American Economic Review* 96, no. 1 (2006): 152-175.
- Papageorge, Nicholas W., and Kevin Thom. "Genes, education, and labor market outcomes: evidence from the health and retirement study." *Journal of the European Economic Association* 18, no. 3 (2020): 1351-1399.

- Pekkarinen, Tuomas, Roope Uusitalo, and Sari Kerr. "School tracking and intergenerational income mobility: Evidence from the Finnish comprehensive school reform." *Journal of Public Economics* 93, no. 7-8 (2009): 965-973.
- Persico, Nicola, Andrew Postlewaite, and Dan Silverman. "The effect of adolescent experience on labor market outcomes: The case of height." *Journal of Political Economy* 112, no. 5 (2004): 1019-1053.
- Price, Alkes L., Nick J. Patterson, Robert M. Plenge, Michael E. Weinblatt, Nancy A. Shadick, and David Reich. "Principal components analysis corrects for stratification in genome-wide association studies." *Nature genetics* 38, no. 8 (2006): 904-909.
- Rietveld, Cornelius A., Sarah E. Medland, Jaime Derringer, Jian Yang, Tõnu Esko, Nicolas W. Martin, Harm-Jan Westra et al. "GWAS of 126,559 individuals identifies genetic variants associated with educational attainment." *science* 340, no. 6139 (2013): 1467-1471.
- Sacerdote, Bruce. "Nature and Nurture Effects on Children's Outcomes: What Have We Learned from Studies of Twins and Adoptees?." In *Handbook of social economics*, vol. 1, pp. 1-30. North-Holland, 2011.
- Schmitz, Lauren, and Dalton Conley. "The long-term consequences of Vietnam-era conscription and genotype on smoking behavior and health." *Behavior genetics* 46, no. 1 (2016): 43-58.
- Schmitz, Lauren, and Dalton Conley. "Modeling gene-environment interactions with quasi-natural experiments." *Journal of personality* 85, no. 1 (2017): 10-21.
- Shanahan, Michael J., Stephen Vaisey, Lance D. Erickson, and Andrew Smolen. "Environmental contingencies and genetic propensities: Social capital, educational continuation, and dopamine receptor gene DRD2." *American Journal of Sociology* 114, no. S1 (2008): S260-S286.
- Solon, Gary. "Intergenerational mobility in the labor market." In *Handbook of labor economics*, vol. 3, pp. 1761-1800. Elsevier, 1999.
- Solon, Gary. "A model of intergenerational mobility variation over time and place." *Generational income mobility in North America and Europe* 2 (2004): 38-47.
- Steelman, Lala Carr, Brian Powell, Regina Werum, and Scott Carter. "Reconsidering the effects of sibling configuration: Recent advances and challenges." *Annual review of sociology* 28, no. 1 (2002): 243-269.
- Sudlow, Cathie, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey et al. "UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age." *Plos med* 12, no. 3 (2015): e1001779.
- Urzua, Sergio. "The effects of cognitive and noncognitive skills on racial and ethnic wage gaps." *Unpublished manuscript, Department of Economics, University of Chicago* (2006).

Vilhjálmsón, Bjarni J., Jian Yang, Hilary K. Finucane, Alexander Gusev, Sara Lindström, Stephan Ripke, Giulio Genovese et al. "Modeling linkage disequilibrium increases accuracy of polygenic risk scores." *The american journal of human genetics* 97, no. 4 (2015): 576-592.

Vytlacil, Edward. "Independence, monotonicity, and latent index models: An equivalence result." *Econometrica* 70, no. 1 (2002): 331-341.

Welland, Jon D. "Cognitive abilities, schooling and earnings: the question of functional form." *The Review of Economics and Statistics* (1978): 622-627.

Young, Alexander I., Seyed Moeen Nehzati, Chanwook Lee, Stefania Benonisdottir, David Cesarini, Daniel J. Benjamin, Patrick Turley, and Augustine Kong. "Mendelian imputation of parental genotypes for genome-wide estimation of direct and indirect genetic effects." *BioRxiv* (2020).