

Intergenerational Mobility and the Timing of Parental Income

Pedro Carneiro

University College London, Human Capital and Economic Opportunity Global Working Group, Institute for Fiscal Studies, Centre for Microdata Methods and Practice, and Centre for Experimental Research on Fairness, Inequality and Rationality–Norwegian School of Economics

Italo López García

RAND Corporation and Institute for Labor Economics

Kjell G. Salvanes

Norwegian School of Economics, Human Capital and Economic Opportunity Global Working Group, and Institute for Labor Economics

Emma Tominey

University of York, Human Capital and Economic Opportunity Global Working Group, and Institute for Labor Economics

We extend the standard intergenerational mobility literature by examining the relationship between adult outcomes of children and the

We thank participants in several seminars for useful comments. Carneiro gratefully acknowledges financial support from the Economic and Social Research Council (ESRC) for the ESRC Centre for Microdata Methods and Practice, through grant ES/P008909/1, and the support of the European Research Council, through grant ERC-2015-CoG-682349. Salvanes thanks the Research Council of Norway for financial support through its Centre of Excellence Scheme FAIR (Centre for Experimental Research on Fairness, Inequality and Rationality) project no. 262675. Tominey acknowledges the Research Council for Norway; the Leibniz Association, Bonn, in the research network “Non-Cognitive Skills: Acquisition and Economic Consequences”; and the British Academy for support. Supplementary material is provided online.

Electronically published February 4, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 3]

© 2021 by The University of Chicago. All rights reserved. 0022-3808/2021/12903-0003\$10.00

timing of parental income during their childhood years, using data from Norway. We find first that, conditional on permanent household income, the child's human capital is higher in households where income is balanced between the early childhood and late childhood years than in households with a more imbalanced income profile. Second, compared to that in the early and late periods of childhood, income in the middle period has relatively low productivity.

I. Introduction

There is a large empirical literature examining the intergenerational transmission of economic status (for recent surveys, see Solon 1999; Björklund and Salvanes 2011; Black and Devereux 2011). It is possible to find estimates of intergenerational mobility for various outcomes and for virtually every country in the world where data linking parents and children are available. Most estimates come from simple models relating a measure of child income and a measure of parental income, such as

$$y_c = \alpha + \beta y_p + u, \quad (1)$$

where y_c is a measure of the child's income, y_p is a measure of parental income, and u is a residual.

Standard economic models of intergenerational transmission justify the use of equation (1) (e.g., Becker and Tomes 1979, 1986), but they collapse the childhood years to a single period of life. More realistic models of parental investments in children distinguish several stages of childhood (e.g., Cunha and Heckman 2007; Cunha, Heckman, and Schennach 2010; Caucutt and Lochner 2012; Cunha 2013). They make clear the idea that the whole history (in particular, the timing) of parental investments in children may be as important as or more important than the total amount invested during the childhood years. In other words, the timing of income may be at least as important as a single measure of income, and the model of equation (1) is missing this issue.

This paper extends the literature on intergenerational transmission by examining the relationship between adult outcomes of children and the timing of parental income during their childhood years, using administrative data from Norway. We focus our attention on two child outcomes: years of schooling and earnings at age 30 (although we also present results for drop-out from high school, college attendance, IQ, and teen pregnancy). To simplify our empirical procedure, we divide childhood in three periods: early (ages 0–5), middle (ages 6–11), and late (ages 12–17).¹ We find that the timing of income matters, over and above permanent income.

Our empirical procedure compares outcomes of children across families with different income profiles across the three periods of childhood

¹ Our results are robust to dividing childhood into more periods.

but the same permanent income.² Although our comparisons are across families, it is easier to describe our results in terms of a single family facing different counterfactual income profiles.

Our measure of permanent income is a statistical one, defined as the sum of (deflated) family income observed during each of the first 17 years of life of the child, discounted to the time of the child's birth (we also consider other measures that consider the whole lifetime of children and parents).³ As shown in Cunha and Heckman (2007) and Caucutt and Lochner (2012), with perfect credit markets and no precautionary savings the timing of income should not matter. In contrast, the timing of income should be important if families face restrictions on borrowing or saving. This is the main rationale behind our empirical procedure.

We show that, keeping permanent income and income in late childhood fixed, child's schooling and earnings improve if a family experiences relatively high (counterfactual) income in early childhood and low income in middle childhood, rather than a relatively low income in early childhood and high income in middle childhood. Stating it more simply, child outcomes improve as (counterfactual) family income rises in early childhood and falls in middle childhood. We also show that a child's schooling and earnings increase when her family income falls in the middle years and increases toward the later years of childhood. Compared to that in the early and late periods of childhood, income in the middle period seems to have low productivity.

Finally, children in families with a balanced profile of income between early and late childhood, keeping fixed permanent income, experience better adult outcomes than those in families with income mainly "front-loaded" in the early period and those in families with income mainly "back-loaded" in the later period of childhood. In other words, we find that there is an inverse-U-shaped relationship between income in the early period relative to income in the late period (which means keeping fixed permanent income and income in middle childhood) and the child's education and labor market outcomes in the adult years.

These empirical patterns are consistent with a model with income uncertainty, partial insurance and imperfect credit markets, and complementarity between investments in children across periods (Cunha and Heckman

² When we represent our results, we keep fixed both permanent income and income in one of the three periods of childhood, e.g., the late period. Then we can compare child outcomes for children growing in families where, e.g., incomes are (relatively) high in early childhood and low in middle childhood with child outcomes for children experiencing low early-childhood income and high middle-childhood income. Note that, conditional on permanent income during the childhood years and income in one period of childhood, incomes in the other two periods of childhood are linearly dependent: when one rises, the other must fall.

³ In other words, this is not a measure of expected future family income at birth (or in any other year) or even a measure of assets, on which parents base their consumption and investment decisions. It is a purely statistical measure used to construct a test of whether the timing of income matters.

2007; Blundell, Pistaferri, and Preston 2008; Cunha, Heckman, and Schennach 2010; Caucutt and Lochner 2012). Income uncertainty, together with partial insurance possibilities such as an inability to borrow and save freely (e.g., for precautionary motives), leads to a setting where investments in children react to parental income shocks in each period. Complementarity between investments taking place in different periods means that human capital is maximized when there is a balanced flow of investments.

In particular, this combination of credit market and technological factors can result in a model where a balanced flow of income shocks may lead to higher human capital than an unbalanced history of shocks. This would be consistent with an inverse-U-shaped relationship between the education of the child and the amount of income that is front-loaded (or back-loaded) in the initial (late) period of childhood. However, our results also suggest that income in the middle period of childhood may be relatively unproductive, compared to income in earlier and later periods.

As mentioned above, we expand equation (1) by replacing the single regressor y_p with multiple measures of income, measured at different stages of childhood. In order to do this, we need data on the entire family income history for a large number of children, which are rarely available. We use data from Norwegian registries for children born during the 1970s, allowing us to link an individual's outcomes as a young adult to the whole history of parental income during the childhood and adolescent years.

We briefly discuss the extent to which we can interpret the association between the child's education and the timing of parental income as causal. We face three (related) challenges. First, the precipitant of the income shock, not the income shock itself, could be the driver of child outcomes. Second, the timing of income could be a choice variable, potentially correlated with investments in children. For example, there are periods when parents take time off work to take care of their children, such as in the case of maternity leave. Third, there is heterogeneity across parents, which can be simultaneously related to income profiles and investments in children. For example, parents on steep income profiles may be more able than those with flatter income profiles. Our discussion is supported by an extensive analysis of these issues, which is included in the appendix (available online).

There exists some literature explicitly examining the role of the timing of income in the formation of human capital. Some of these papers use survey data from the United States and Germany and rely on relatively small data sets (Duncan et al. 1998; Levy and Duncan 2000; Jenkins and Schluter 2002; Carneiro and Heckman 2003; Caucutt and Lochner 2012). Others use much larger register data sets for Denmark and Norway (Aakvik, Salvanes, and Vaage 2005; Humlum 2011), but nevertheless they estimate very restrictive models. In particular, all these papers estimate

regressions of child outcomes on the income of parents at different ages. Since the levels of income in different periods enter in a linear and additive way in these models, they are assumed to be “substitutes” in the production of human capital.

Relative to the papers using US and German survey data, our paper relies on much better data (larger samples and richer income histories), allowing us to estimate much more flexible models with considerable precision. This is true even when we contrast our analysis to the ones using register data for Norway and Denmark. The flexible models we estimate enable us to construct a richer picture of the role of the timing of income than the one presented in previous work. This is important because the complementarity (or other interactions) of investments in human capital across periods (Cunha, Heckman, and Schennach 2010) may translate into complementarity (or other interactions) of income shocks across periods.⁴

In another related paper, Caucutt and Lochner (2012) present an overlapping-generations model of parental investments in children. In addition, Cunha (2013) also presents such a model with overlapping generations and multiple periods of parental investments, explicitly accounting for uncertainty in income. The model we consider (and discuss in detail in the appendix) is much simpler than the models in these two papers. It is neither a general equilibrium model nor an overlapping-generations model, but it serves much more modest purposes, including (1) to examine whether it is possible to explain the main patterns in the data using a parsimonious economic model of parental investments and (2) to conduct simple simulations of the how the insurance provided by the tax and benefit system can influence the effect of the timing of family income on child outcomes.

The structure of the paper is as follows. In section II we discuss the specifics of the Norwegian welfare state, in section III we describe the data, and in section IV we present our empirical methods. Section V discusses our empirical results. Section VI provides some interpretation of our results, including a brief discussion of problems caused by the potential endogeneity of the timing of income, the role of insurance possibilities, and a simulation of simple dynamic models of parental investments in children. Finally, section VII concludes.

II. The Norwegian Welfare State

The Norwegian welfare state—together with the welfare states in Sweden and Denmark—has social insurance and economic equality as core elements.

⁴ Note that in a complementary paper, Hilger (2016) exploits parental layoffs to estimate the effect of parental job loss when children are aged 12–29 on children’s long-term outcomes, including college enrollment and earnings to age 25, but does not distinguish between the effects of different timings of this shock.

The Scandinavian model has been slowly developed over time, where the most important developments took place between World War II and the 1990s. For our purpose of understanding the effect of income shocks on children and teenagers, we provide an overview of important aspects of the social security and transfer system and family policy.

Our sample period starts with births in 1970 and measures the latest child outcomes in 2014. Across this period, the Norwegian welfare state has consistently been characterized by highly progressive taxation, an extensive transfer system to redistribute income, and a strong social insurance system. Social security is characterized by universal benefits that cover the whole population rather than means testing, with many benefits and transfers related to current or previous labor force participation (Lindbeck 1997; Barth, Moene, and Willumsen 2014).

The social security system, disability insurance, and sick leave programs are all based on income replacement, whereby a worker is entitled to a benefit of a fixed percentage of the previous year's pay or of the average over the past three years. (See Markussen and Røed 2014 and Autor et al. 2016 for details on disability insurance and sick leave programs).

A policy area that has seen reform in our sample period is family and active labor market policies, including maternity leave and childcare. The first maternity leave reform took place in 1977, providing full income replacement for working mothers for four-and-a-half months (Carneiro, Løken, and Salvanes 2015). This reform was extended successively in the 1980s and 1990s up until today, with access to 49 weeks with 100% coverage (Dahl et al. 2016). Universal and strongly subsidized childcare was established in the late 1970s for children aged 1–5 years but was really expanded in the 1980s and 1990s (Havnes and Mogstad 2011; Black et al. 2014). Therefore, this reform is unlikely to have affected children from our sample, who were born in the period 1971–80.

In addition to the family policies, an important area of change in our data period comes from changes in the progressivity of the taxation of income. The tax system has become less progressive over time through changes in the level of deductions and especially of surtaxes (Thoresen 2004). The major changes took place in 1992, 1995, 1999, 2001, 2004, and 2006 (Fjærli and Aaberge 1999; Blundell, Graber, and Mogstad 2015). One example is the substantial tax reform of 1992, when personal income taxation was changed from a system where the nominal tax rate on income varied from 26.5% to more than 40% to a dual system where there was still progressive taxation of labor income but capital income was taxed at a flat nominal rate of 28%. Moreover, the tax base was increased such that there was less room for deductions, such as those on mortgage interest payments. There were several other changes as well as subsequent tax reforms, which together led to a reduction in the progressivity of the system

as a whole. As documented in Blundell, Graber, and Mogstad (2015), the average tax rates were also strongly reduced across the period, with the consequence of a reduction in progressivity. Many of these changes were related to similar changes in all OECD (Organisation for Economic Co-operation and Development) countries in the same period (OECD 1998).

III. Data

Our data source is the Norwegian Registry, maintained by Statistics Norway, for the years ranging from 1971 to 2014. It is a linked administrative data set that covers the population of Norway, and it is a collection of different administrative registers providing information about (among other things) month and year of birth, educational attainment, labor market status, earnings, and a set of demographic variables (age, gender) as well as information on families, including parents' marital status. We are able to link individuals to their parents, and it is possible to gather labor market information for both fathers and mothers.

For the bulk of our analysis we select all births in the period 1971–80. In particular, for our main analysis we construct annual household net earnings data for each year, from the three years preceding the child's birth through to their twentieth birthday (we also construct earnings for the entire life cycle of parents, which are used in our sensitivity analyses). The net income measure is calculated net of a progressive income tax, cash transfers, child benefits (for children up to age 18), unemployment and sickness benefits, general deductions for expenditures going to work, and a regional compensation for living in the northernmost region. For some sensitivity analyses we analyze the effect of the timing of gross income, which is measured as pensionable income, taking out sickness and unemployment benefits. This gives us information on income data and parental characteristics, mapped to children's outcomes, for 499,442 children.

We deflate income to 2000 and discount all incomes to the year of birth of the child, using a real interest rate of 4.26% (Aakvik, Salvanes, and Vaage 2005). However, our results are robust to a large range of fixed discount rates between 0% and 15% and to time-varying discount rates (matching the risk-free interest rate evolution during the relevant years for our analysis). In order to construct a measure of income in each of the three periods, we take the sum of discounted annual household incomes within each period (0–5, 6–11, 12–17). Permanent income is then defined as the sum of discounted income in the three periods.

Our data include a large range of human capital outcomes, including years of education for each individual, an indicator for dropping out of high school at the age of 16, and college enrollment. The consequence

of early drop-out is that individuals do not receive a certificate for vocational or academic achievement, which severely limits opportunities in the labor market and prohibits access to further education. Unfortunately, it is not possible to measure whether a college degree was completed. Individuals are at least 26 years of age when we observe their educational achievement, and consequently, most of them can be expected to have exited school.

We also present results in the appendix for other outcomes, including earnings at age 30, IQ, and teenage pregnancy. Military service is compulsory in Norway for males, and between the ages of 18 and 20 males usually take an IQ test. This test is a composite of arithmetic, word,⁵ and figure tests,⁶ all of which are recognized as tests of IQ. For teen pregnancy, we construct an indicator for whether an individual (female) had a child when she was between the ages of 16 and 20.

Finally, we use a detailed set of control variables, which are important to ensure that we are comparing differing income profiles across otherwise similar households. First, we construct household-specific slopes of income profiles by taking the difference between household income when the child is aged 18–20 and that in the three years before birth.⁷ Other controls include the child's year of birth and gender as well as parental years of education and age at birth. Family size and marital status are endogenous and so are not included in our benchmark specification but rather are included as a robustness check.

The descriptive statistics for the sample are reported in table 1. There are 499,442 child-level observations, which include all individuals born in Norway between 1971 and 1980 for whom we were able to collect household income data. The permanent income of the household (in the period between the ages of 0 and 17 of the child) is about £312,200. There is substantial income dispersion: the standard deviation is £93,500. Income in each period (1, 2, and 3 for early, middle, and late childhood) falls with the age of the child because of discounting.

Mothers have, on average, 11.15 years of schooling, which is slightly lower than the average years of education of the fathers (11.47). Mothers are much younger than fathers at birth (26 vs. 29 years of age). The average years of education of the children in our sample is 12.74. Twenty-one percent of children drop out of high school, but 40% attend college. The average annual earnings of these children at age 30 is £21,376. Regarding additional child outcomes, IQ is available for only males and takes values on a nine-point scale, with a sample average of 5.26 and a

⁵ The word tests are most similar to the Wechsler Adult Intelligent Scale.

⁶ The figure tests are similar to the Raven Progressive matrix.

⁷ A robustness check conditions instead for income growth rather than the level difference between pre- and postchildhood income.

TABLE 1
DESCRIPTIVE STATISTICS

	Observations	Mean	Standard Deviation
Household income:			
Household income, period 1: age 0–5	499,442	10.72	3.30
Household income, period 2: age 6–11	499,442	10.55	3.30
Household income, period 3: age 12–17	499,442	9.95	4.24
Household permanent income: age 0–17	499,442	31.22	9.35
Covariates:			
Mother years of schooling	499,442	11.15	2.71
Father years of schooling	499,442	11.47	3.01
Mother age at birth	499,442	26.25	5.02
Father age at birth	499,442	28.99	5.71
Child year of birth	499,442	1,975.27	2.89
Child outcomes:			
Years of schooling	499,442	12.74	2.41
High school dropout	499,442	.21	.41
College attendance	499,442	.40	.49
Log earnings at age 30	295,380	9.97	.79
IQ (males only)	238,184	5.26	1.79
Teenage pregnancy (females only)	238,611	.08	.28

NOTE.—Income and earnings variables are in £10,000s, 2000 prices.

standard deviation of 1.79. Finally, teen pregnancies occur for 8% of the females in our sample.⁸

IV. Methods

Let Y_i be an outcome for child i (education, high school completion, college attendance, earnings, IQ, teenage pregnancy) measured in late adolescence or young adulthood. We are interested in Y_i as a function of the history of (discounted) household income I_{it} in each period t ($t = 1, 2, 3$) and the permanent income of the parents, PI_i . Since $PI_i = I_{i1} + I_{i2} + I_{i3}$, we drop one of the periods from the model, say I_{i1} . Therefore, we write

$$Y_i = m(PI_i, I_{i2}, I_{i3}) + Z_i\delta + \varepsilon_i. \quad (2)$$

The term Z_i denotes a large set of controls, which aim to capture that parents facing different income profiles may also be different in other dimensions. We condition on paternal and maternal education interacted with paternal or maternal age at birth (by including dummies for years of education and age at birth for each parent interacted with each other). This controls for different age-education profiles across households. Moreover, we construct a measure of household income growth between

⁸ The income process is studied in detail in Carneiro, Salvanes, and Tominey (2015). They find that, as in many other countries, the income process for Norwegian fathers can be fairly well approximated by the sum of a random walk (permanent shock) and a low-order MA (moving-average) process (temporary shock).

the ages of 0 and 17 of the child, based on income 1–3 years before birth (prebirth income) and income 1–3 years after age 17 (post-17 income). This means that we explore fluctuations in income around deterministic age-income profiles that are allowed to vary with education, after accounting for heterogeneous income growth (and, of course, keeping fixed permanent income). The remaining controls in the model are the birth year and gender of the child.

The term ε_i should be interpreted as the unobserved heterogeneity that is left after we control for permanent income and controls. We assume that ε_i has finite conditional variance, $E(\varepsilon_i^2 | \text{PI}_i, I_{i2}, I_{i3}, Z_i) \leq C < \infty$, and that $E(\varepsilon_i | \text{PI}_i, I_{i2}, I_{i3}, Z_i) = 0$. We are interested not in the impact of PI itself on Y but on the impact of the timing of income (I_2 and I_3) on Y , after conditioning on PI and Z . In other words, we want to compare outcomes of children whose parents have the same level of permanent income between the ages of 0 and 17 and the same age-education profiles, household income growth, and child's birth year and gender but differ in the level of income they experience in each period.

The assumption that $E(\varepsilon_i | \text{PI}_i, I_{i2}, I_{i3}, Z_i) = 0$ may be controversial. We would like to interpret I_{i2} and I_{i3} as income shocks orthogonal to other (unobservable) determinants of outcomes Y_i , conditional on PI_i and Z_i . It is likely that PI and Z_i absorb much of the relevant unobserved heterogeneity across parents (correlated with the overall level of their income). We show in section VI.A and appendix A3 that when we explore plausibly exogenous changes in the timing of income across households, resulting from changes in the tax and benefit system, our results are unchanged.⁹

We allow $m(\text{PI}_i, I_{i2}, I_{i3})$ to be a nonparametric function of its arguments. It tells us what counterfactual (adult) outcome Y_i would be if a household faced a particular counterfactual income profile, $m(\text{PI}_i, I_{i2}, I_{i3})$. It is important to be able to estimate a flexible function, and the reason is the following. Parents are faced with income shocks in each period and, in response, decide how much to invest in children (and how much to consume and save). There is a technology that links the adult human capital of an individual to the whole history of parental investments in childhood and adolescence. The link between income shocks and child outcomes, described by equation (2), depends on many factors, including preferences,

⁹ We also present results where we control for marital breakup and number of children. In sec. VI.A, we repeat several of these arguments and provide a more detailed discussion of potential violations of the assumption that $E(\varepsilon_i | \text{PI}_i, I_{i2}, I_{i3}, Z_i) = 0$. In addition, one implication of our assumption is that prebirth investments should be uncorrelated with the timing of income but may still affect child outcomes. We test this in app. A3 by examining the relationship between having a low-birth-weight baby and the subsequent timing of parental income. We show that, although low birth weight is strongly correlated with PI_i , it is uncorrelated with I_{i2} and I_{i3} . Similarly, we show that income shocks occurring after the period in which child outcomes are measured do not drive those outcomes.

technology, information, and the structure of credit markets (insurance possibilities). Therefore, there are likely to be complex nonlinearities and interactions between the different income measures included in the model.

We are particularly interested in $m_2(\text{PI}_i, I_{i2}, I_{i3}) = \partial m(\text{PI}_i, I_{i2}, I_{i3}) / \partial I_{i2}$ and $m_3(\text{PI}_i, I_{i2}, I_{i3}) = \partial m(\text{PI}_i, I_{i2}, I_{i3}) / \partial I_{i3}$. The term $m_2(\text{PI}_i, I_{i2}, I_{i3})$ tells us how the counterfactual outcome Y_i for a child in a particular family would change (starting from some benchmark) if we reduced counterfactual income in period 1 and raised it in period 2, keeping PI_i and I_{i3} fixed (and $\text{PI}_i = I_{i1} + I_{i2} + I_{i3}$). In section V, we present a series of graphs relating Y_i and I_{i2} (for different outcomes Y). The graphs vary, depending on the values of PI_i and I_{i3} on which we evaluate this function. All other controls Z_i are held at their mean values. An analogous interpretation and graphical representations of results can be given to $m_3(\text{PI}_i, I_{i2}, I_{i3})$.

Details of the nonparametric estimation methods are in appendix A1.

V. Results

In this section, we present semiparametric estimates of $m(\text{PI}, I_2, I_3)$, following the method described in appendix A1.¹⁰ The empirical results in this paper are presented through a series of graphs. We first focus on years of education as the outcome of interest and then discuss results for earnings at age 30, IQ, and teenage pregnancy. The outcomes high school drop-out and college attendance are mentioned in the text, but results are presented only in the appendix. In order to implement the estimator, we first need to create a grid of evaluation points for $m(\text{PI}, I_2, I_3)$. We take 19 points for each income variable (PI, I_2, I_3), corresponding to the ventiles of each variable's distribution. This gives us a tridimensional grid with 6,859 points ($19 \times 19 \times 19$).

It is standard practice to trim the data so as to avoid spurious results driven by small cells. Therefore, we drop 2% of the sample, corresponding to the cells with the smallest number of observations. In our main results, we use a uniform kernel and choose the bandwidths using the formula in equation (1) of appendix A1, setting $C = 6$. Below, we show that our results are robust to the choice of bandwidths.

One way to present our estimates of $m(\text{PI}, I_2, I_3)$ is through a series of two-dimensional graphs, where in the y -axis we represent the outcome of interest and in the x -axis we represent one of the income variables in $m(\text{PI}, I_2, I_3)$. The advantage of this presentation is that the graphs are straightforward to read. The downside of this type of presentation is that it allows variation of only one income period at a time, which means that we need to fix the remaining two income variables at some predetermined

¹⁰ Appendix A2 estimates the effect of the timing of income in a parametric setting, with corresponding figs. A1 and A2.

values (we fix the remaining control variables at their mean values). Therefore, we create multiple curves for each outcome, corresponding to different fixed values for the left-out income variables in each graph.

For each outcome, we present three graphs. In the first graph, we fix PI and I_3 at three different values each (the third, fifth, and seventh deciles of the distribution of each variable) and vary only I_2 , for a total of nine possible combinations. These are presented in one figure, which plots $m(\text{PI}, I_2, I_3)$ against I_2 (for given values of PI and I_3).

Since $\text{PI} = I_1 + I_2 + I_3$, when PI and I_3 are fixed at particular values, we cannot vary I_1 and I_2 independently. Therefore, by moving toward the right along the x -axis in each graph, we are able to see how the outcome varies as I_1 decreases and, simultaneously, I_2 increases (i.e., as income is reduced in period 1 and increased in period 2). In other words, our estimates concern the impact of changes in counterfactual family income profiles across different pairs of periods, holding constant permanent income and one per-period income (I_3 in this particular case). The support of I_2 over which we can evaluate $m(\text{PI}, I_2, I_3)$ is not the same across all graphs within a figure because, for different combinations of PI and I_3 , there are values of I_2 for which there are no observations in the sample.¹¹ The second graph keeps PI and I_2 fixed and lets I_3 vary (so that income is reduced in period 1 and increased in period 3). The third graph keeps PI and I_1 fixed and varies I_3 (a reduction in period 2 income and simultaneous increase in period 3 income).

In order to understand whether the patterns we uncover are statistically meaningful, in each graph we display two other parameters, α_1 and α_2 (and respective standard errors), which are defined as follows. For each panel, let H be the highest grid point for the income variable being used in that panel, let L be the lowest grid point, and let M be the median grid point (corresponding to the 50th percentile of the distribution of that income variable in that panel). Take the case where we fix $\text{PI} = \overline{\text{PI}}$ and $I_3 = \overline{I_3}$ and we let I_2 vary. Then, define

$$\begin{aligned}\alpha_1 &= m(\overline{\text{PI}}, M, \overline{I_3}) - m(\overline{\text{PI}}, L, \overline{I_3}), \\ \alpha_2 &= m(\overline{\text{PI}}, H, \overline{I_3}) - m(\overline{\text{PI}}, M, \overline{I_3}).\end{aligned}\tag{3}$$

The value of α_1 is the difference between the values the outcome takes in the median and the lower extreme of the support of I_2 , while α_2 is the difference between the values the outcome takes in the median and the upper extreme of the support of I_2 . If $m(\overline{\text{PI}}, I_2, \overline{I_3})$ did not vary with I_2 (in which case the timing of income would be irrelevant, at least when we

¹¹ Note that we have data on the universe.

compare first- and second-period incomes), we would expect $\alpha_1 = \alpha_2 = 0$, so these parameters help us quantify the importance of the timing of income. Although there are several other ways to assess whether the curves we estimate are statistically different from what we would obtain if the timing of income was irrelevant (a flat line), the one we present here is particularly simple to implement and understand.

A. *Years of Schooling*

We start by reviewing the evidence of the effect of the timing of parental income on years of schooling of the child. In figure 1A, we represent the relationship between years of schooling of the child (Y_i) and I_2 , keeping fixed PI and I_3 (which means that I_1 decreases when I_2 increases, so we are measuring the impact on years of schooling of a counterfactual experiment that raises I_2 while simultaneously decreasing I_1). Figure 1B represents the relationship between Y_i and I_3 , keeping fixed PI and I_2 (impact of an increase in I_3 and simultaneous decrease in I_1), and figure 1C represents the relationship between Y_i and I_3 , keeping fixed PI and I_1 (impact of an increase [decrease] in I_3 [I_2]).

Each panel represents nine functions, which correspond to nine different points of evaluation for two of the income variables omitted from the figure. For example, in A we keep PI and I_3 fixed (while I_2 , and therefore I_1 , are allowed to change) at the third (low), fifth (medium), or seventh (high) deciles of the distributions of each of these variables (since there are two omitted variables and three evaluation points for each of them, we get a total of $3^2 = 9$ curves). This allows us to see how the relationship between Y_i and I_2 varies when we change the values of PI and I_3 .

It is useful to discuss one curve in one panel in detail. Then one can easily understand how to read all curves across all panels. Take the solid line in figure 1A, corresponding to low PI and low I_3 . At the leftmost point of this line, corresponding approximately to $I_2 = \text{£}60,000$, the average child has about 12.5 years of schooling. Since $\text{PI}^{\text{low}} = \text{£}250,930$ and $I_3^{\text{low}} = \text{£}81,000$, if $I_2 = \text{£}60,000$, then $I_1 = \text{£}250,930 - \text{£}60,000 - \text{£}81,000 = \text{£}109,930$. As I_2 increases from $\text{£}60,000$ to $\text{£}120,000$ or so (and I_1 decreases from $\text{£}109,930$ to $\text{£}49,930$), Y_i falls by about half a year of schooling, from about 12.5 to below 12.

Within each panel of figure 1, even though different curves are evaluated at dramatically different values of the omitted income variables and clearly have different locations, they display strikingly similar shapes. These similarities can be better seen in figure 2, where we adjust all the curves in each panel to have a similar location. More precisely, in figure 2 the graphs have been adjusted along the y-axis to have a common mean value of the child outcome, and in a small number of curves there was also an adjustment along the x-axis so that the peak of each curve occurs at a

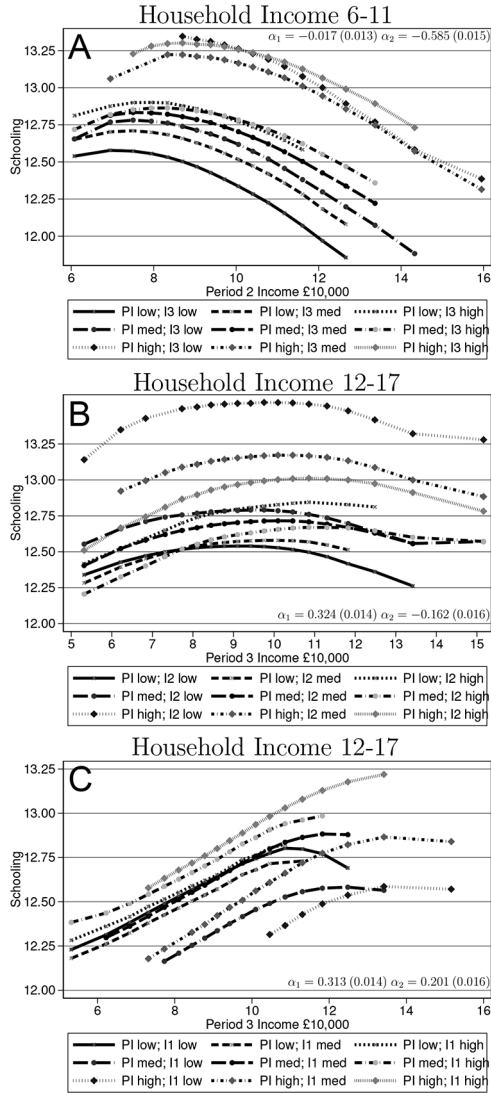


FIG. 1.—Semiparametric estimates; dependent variable is years of schooling; income in 2000 prices, £10,000s. Semiparametric estimates control for dummies for paternal education interacted with age and maternal education interacted with age, household net income profile, gender, and child year of birth, and they cluster at the family level. α_1 and α_2 and associated standard errors, defined in equation (3), indicate the significance of the curvature, denoting differences in outcome between the median and the lower extreme and between the upper extreme and the median, respectively. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively.

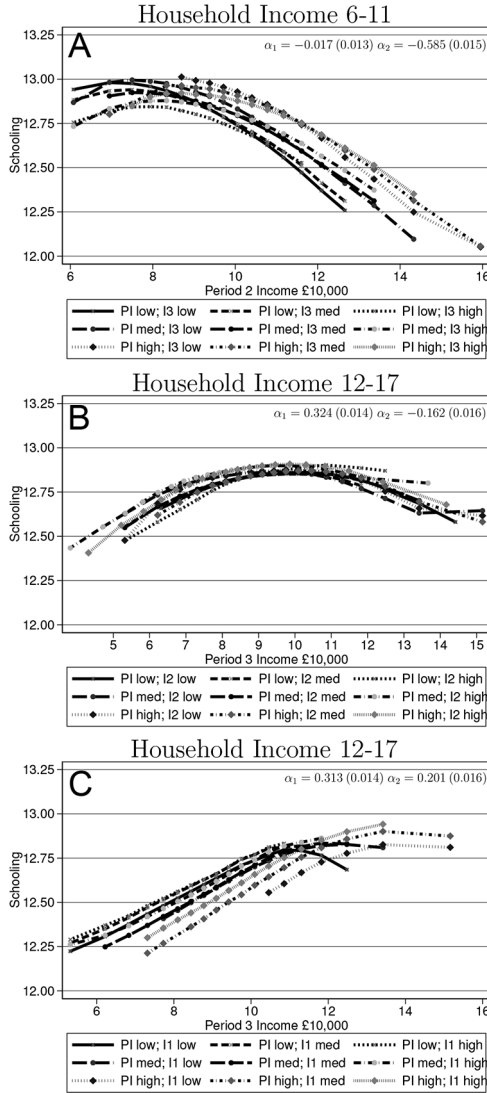


FIG. 2.—Semiparametric estimates; dependent variable is years of schooling; adjusted figures; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x-axis and, for B, the y-axis are described in sec. V.A.

similar value of x .¹² The adjustments are additive (adding constants horizontally and vertically to the curves in fig. 1), so they do not influence the curvature of the functions shown here. In the rest of the paper, we present graphs with adjustments similar to those in the panels of figure 2 (the appendix shows the full set of figures before adjustment of the axes).

To test whether the patterns in the figures are statistically meaningful, we compute α_1 and α_2 (and report standard errors for the test that these parameters are equal to zero). There is one α_1 and one α_2 for each curve in each panel of figure 2, so for each panel we report the average values of α_1 and α_2 across the nine different curves and the corresponding standard errors (the individual values of these parameters are shown in fig. A3; figs. A1–A40 are available online).

All of the curves in figure 2A (and fig. A3a–A3c) are quite similar. The dominant pattern is a downward-sloping curve, whereby an increase in I_2 and a decrease in I_1 together lower years of schooling. The slope of this curve is statistically different from zero (especially when we look at α_2). This suggests that income occurring at ages 0–5 has a relatively larger impact on years of schooling than income occurring during ages 6–11.

For very low levels of I_2 (very high levels of I_1), an increase in I_2 initially raises years of schooling in some of these figures. Close inspection of α_1 in figure 2A (and fig. A3a–A3c), however, tells us that this initial increase is not generally statistically different from zero.¹³

The slopes of the curves in figure 2A become steeper as the level of I_1 falls, as evidenced from a comparison of α_1 and α_2 . Here, α_1 takes the value of -0.017 , suggesting that a reduction of I_1 and a simultaneous increase of I_2 , from the bottom to the middle of the distribution of I_2 , lowers schooling by 0.017 years. The corresponding figure for α_2 suggests that a shift from I_1 to I_2 , from the midpoint of I_2 to the highest level, lowers schooling by 0.585 years. This says that the larger impact of income at 0–5, compared to that of income at 6–11, is especially large when I_1 is low (rather than when it is high).

Consider now the productivity of I_1 relative to that of I_3 . In figure 2B (and fig. A3d–A3f), I_2 and PI are fixed within each graph at the third, fifth, and seventh deciles, and we show how the estimate of $m(\overline{PI}, \overline{I_2}, I_3)$

¹² Specifically, for fig. 2B, one curve was shifted by a value of 1 to the right (PI low and I_2 low), and two curves were shifted by the value of 1 to the left (PI medium and I_2 high; PI high and I_2 high).

¹³ As a benchmark, it is useful to compare the downward slopes in the different curves of fig. 2A, with the slope of the relationship between years of schooling and permanent income shown in fig. A1. The slopes of the curves in fig. 2A are of a very similar magnitude to those shown in fig. A1. For example, a £50,000 increase in permanent earnings leads to about an extra 0.5 years of education, while an increase in middle-childhood income from £80,000 to £130,000 leads to a change in schooling of the same magnitude, about 0.50 years of education. It is remarkable that, even conditional on permanent income, the impact of the timing of income is as large as the impact of permanent income.

increases with I_3 (corresponding to decreasing levels of I_1). There is an inverse-U-shaped pattern for nearly all the curves displayed in the figure. Years of schooling initially increase as income is moved from (initially high) I_1 to (initially low) I_3 , reach a maximum at around £100,000 of (discounted) I_3 , and fall thereafter. Here, α_1 is positive and statistically different from zero, and α_2 is negative and statistically different from zero.

What this means is that, across different values of PI and I_2 , the years of schooling of the child are maximized when there is some balance between early- and late-childhood income. If income is too concentrated in either early or late childhood, then one can improve education outcomes of children by “shifting” income toward the other period. Looking at the estimate for α_1 , moving income from a low initial level of I_3 (and a high level of I_1) to the median level of I_3 raises schooling by 0.324 years. Looking at α_2 , moving income from the median level in period 3 toward very high levels of I_3 (and consequently very low levels of I_1) reduces schooling by 0.162 years.

Figure 2C (and fig. A3g–A3j) examines the trade-offs between income in middle and late childhood (keeping fixed income in early childhood and permanent income). The curves display an increasing shape, indicating that years of schooling are maximized when income is reduced in middle childhood and increased in late childhood. Across curves, α_1 and α_2 are both positive and statistically different from zero. The magnitudes of these impacts are substantial and similar, if slightly smaller than comparable changes in permanent income. The slopes of the estimated curves are particularly steep when I_3 is increased from initially low levels, where α_1 and α_2 take the values of 0.313 and 0.201, respectively.

It is striking how similar the displayed patterns are across the different curves within each panel of figures 1 and 2. We have essentially decreasing functions when looking at $m(\overline{PI}, I_2, \overline{I_3})$, inverse-U shapes when looking at $m(\overline{PI}, \overline{I_2}, I_3)$, and increasing functions when displaying $m(\overline{PI}, \overline{I_1}, I_3)$. We do not impose this in our estimation procedure, and we were surprised that we did not have more variety in the estimated patterns, depending on, for example, how wealthy the family was (in terms of permanent income). Our initial goal was to have a framework flexible enough to uncover diversity in the behavior of these functions, but what we uncovered was instead stability. The stability in the shapes of these functions is also observed when we study other outcomes, as we show below.

These results we show have three important implications. First, the timing of income shocks is relevant for human capital formation. If the timing of income was irrelevant, then all these graphs would be horizontal lines, with only permanent income being relevant for human capital outcomes. It is likely that the timing of income shocks affects the timing of investments in human capital (which in turn affects human capital formation). This will happen if parents have imperfect insurance possibilities against income shocks.

Second, if the timing of income matters because parental investments react to income shocks, the shape of the curves in figure 2*B* is consistent with an underlying technology of skill formation that exhibits complementarity in investments across periods. The reason is that, under complementarity, it is desirable to maintain a balanced flow of investments over the life of the child. However, at this point, we cannot distinguish this model from any other model where smoothing investments is desirable (see the discussion in app. A6).

Third, our results are also consistent with a model where investments in the second period of childhood (ages 6–11) have relatively lower productivity than investments in early childhood and late childhood. Given that all our outcomes are measured in late adolescence and adulthood, it is perhaps not surprising that later investments have “higher productivity” than earlier investments, just because they have had less time to depreciate. If anything, what is surprising is that early investments can be so much more important than investments in middle childhood and as important as investments in late childhood. This may happen because they depreciate less than later investments and/or because they are complementary with investments in late childhood.

It is important to understand the upward-sloping section in each curve (see also app. A5). One might think that, for a given level of permanent income, it should not be worse to have all income available in the first period (or in an earlier period) than to receive it in payments spread out over different periods of childhood. This is because, if permanent income becomes fully available at time zero, then one can allocate it freely across periods just by saving the appropriate amount, regardless of whether or not one can borrow. If this is true, then children in households getting all their income in early childhood (and nothing in a later period) should not do worse than children in households with a more spread-apart flow of income (keeping permanent income fixed).

However, the reasoning just described ignores the fact that incomes are uncertain and that parents cannot fully anticipate future shocks. They cannot be certain of whether their income will rise or decline in future periods. In addition, their objective function is not just maximizing the net present value of child outcomes, but they also trade off their own consumption for investments in children (unless capital markets are perfect). As a result, in the presence of income uncertainty, when faced with income shocks parents might oversave for precautionary motives and invest in children less than they could, unless they have perfect insurance. An income shock occurring late in the life cycle can induce additional investments in this period of life that might not have occurred if the shock had taken place earlier in life. These investments could in turn lead to an improvement in child outcomes.

B. *Gender Differences*

There is a large literature documenting gender differences in parental investments in child human capital in developing countries (e.g., Rosenzweig and Schultz 1982; Pitt, Rosenzweig, and Hassan 1990; Lee, Parish, and Willis 1994; Quisumbing 1994). In developed countries, parental investments and the returns to these investments have also been found to be different between boys and girls. Baker and Milligan (2016) document higher levels of parental investments that promote learning in girls than in boys, which lead to higher scores for reading and math achievement of girls, compared to boys. Bertrand and Pan (2013) find that the response of children's non-cognitive skills to parental investments is much more sensitive for boys than for girls, suggesting that the production function that translates investments to child outcomes may vary across child gender.

In figure 3 (and fig. A4), the semiparametric estimates of the effect of the timing of income on years of schooling are shown for girls, while estimates for boys are in figure 4 (and fig. A5). The curves for girls are similar to those reported above, especially in figures 3A and 3C. There is a less pronounced inverse-U-shaped pattern for figure 3B. In contrast, for boys there is a more pronounced inverse-U-shaped pattern in figure 4B (concerning the productivity of I_3 relative to I_1) than when we looked at the whole sample, and several curves in figure 4A (concerning the productivity of I_2 relative to I_1) also display a strong inverse-U shape (confirmed by the estimates of α_1 and α_2 in the same panel). Figure 4C is similar to figures 2 and 3.

As explained above, one possible explanation for inverse-U-shaped patterns in the relative productivity of income in one age versus another is consistent with a model where investments in children react to income shocks and there is dynamic complementarity between investments in different time periods. Our estimates suggest that, for boys, this complementarity is more important than it is for girls.

C. *Other Outcomes: Earnings at Age 30, IQ, Teenage Pregnancy*

In addition to years of schooling, we considered high school dropout rates and college attendance rates separately, since they correspond to education decisions at the lower and upper tails of the education distribution. Perhaps not surprisingly, these results are broadly similar to the ones we showed for years of education, and we display these results in the appendix (figs. A7–A10). More interestingly, we examine what our results look like when we consider three other outcomes for the child: log earnings measured at age 30, IQ, and teenage pregnancy. We turn to these next.

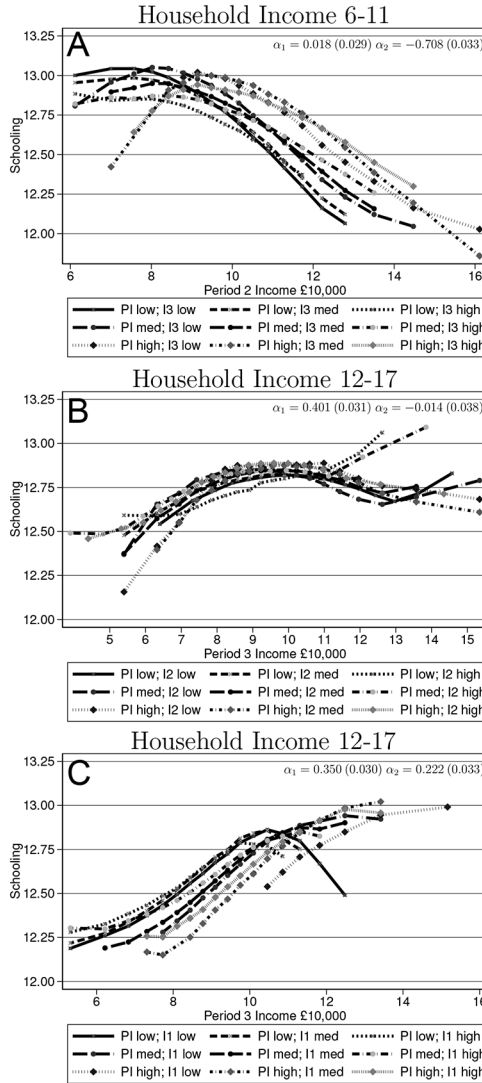


FIG. 3.—Semiparametric estimates; dependent variable is years of schooling; girls; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x-axis and, for B, the y-axis are described in sec. V.A.

Most of the literature on intergenerational mobility links the earnings of parents to the earnings to children. Figure 5 (and fig. A6) shows semiparametric estimates analogous to those discussed above (e.g., fig. 2), but using as the outcome the log annual earnings of children at age 30. Figure 5A

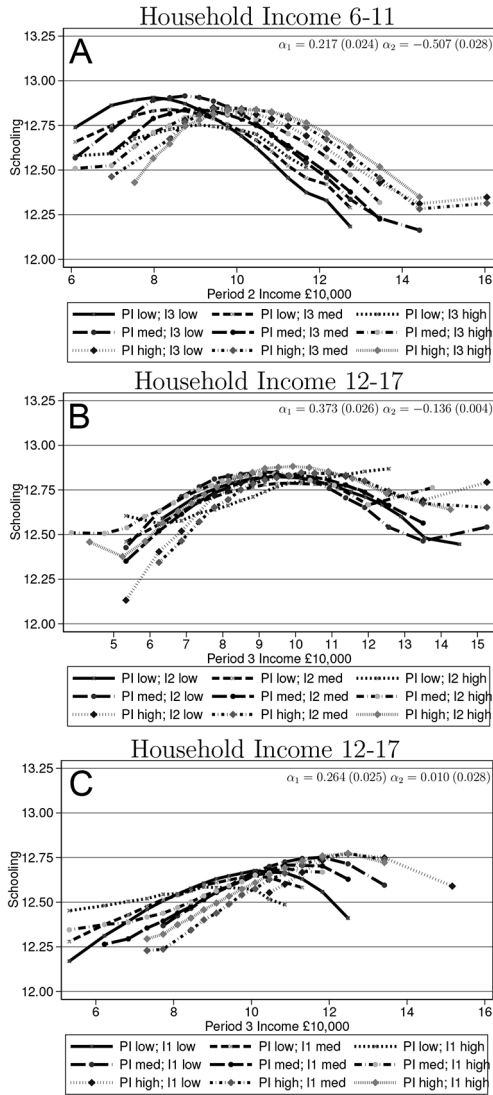


FIG. 4.—Semiparametric estimates; dependent variable is years of schooling; boys; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x-axis and, for B, the y-axis are described in sec. VA.

(fig. A6a–A6c) shows a downward-sloping curve (I_2 vs. I_1), figure 5B (fig. A6d–A6f) shows an inverse-U-shaped curve (I_3 vs. I_1), and figure 5C (fig. A6g–A6j) shows an upward-sloping curve (I_3 vs. I_2). These are remarkably similar to the schooling results, suggesting that whatever

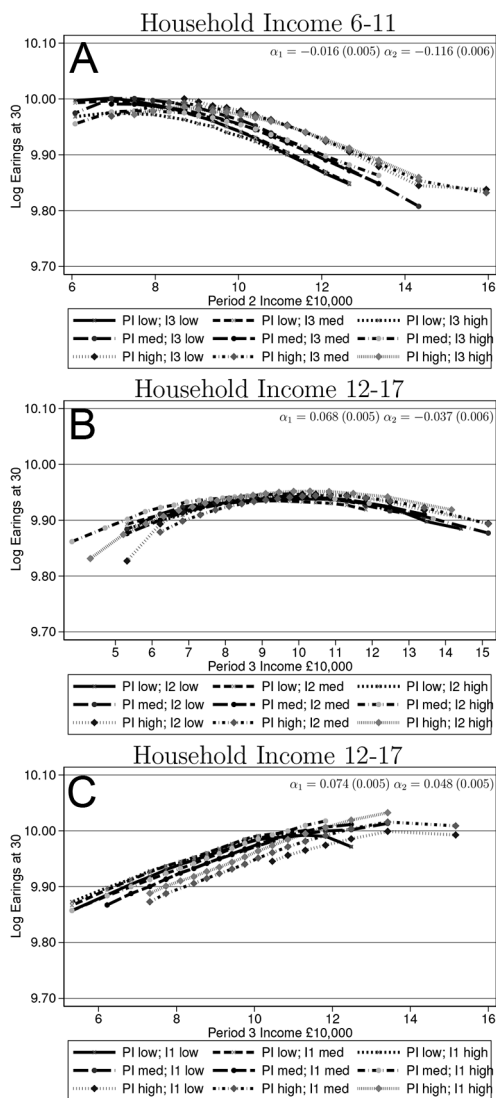


FIG. 5.—Semiparametric estimates; dependent variable is log earnings at 30; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x -axis and, for B , the y -axis are described in sec. V.A.

component of human capital is driving the schooling results may also be driving what we observe for earnings.

Figures 6 and 7 (figs. A11, A12) show our estimates when we use IQ and teenage pregnancy, respectively, as the outcomes of interest. It is important

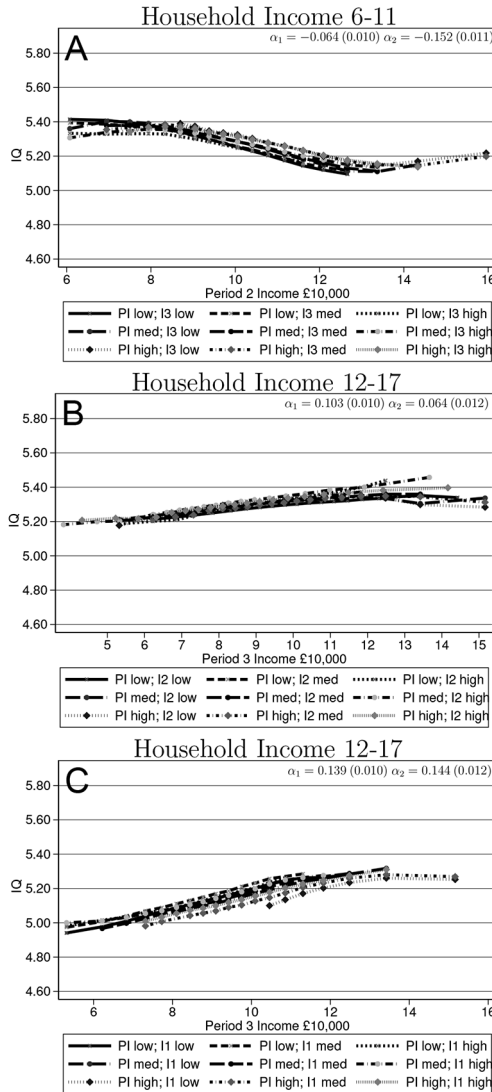


FIG. 6.—Semiparametric estimates; dependent variable is IQ; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x-axis and, for B, the y-axis are described in sec. V.A.

to look at outcomes such as these because they could potentially reflect dimensions of human capital different from those manifested through schooling and earnings. IQ is measured through the armed forces tests at age 18, only for males, and is traditionally seen as a good measure of

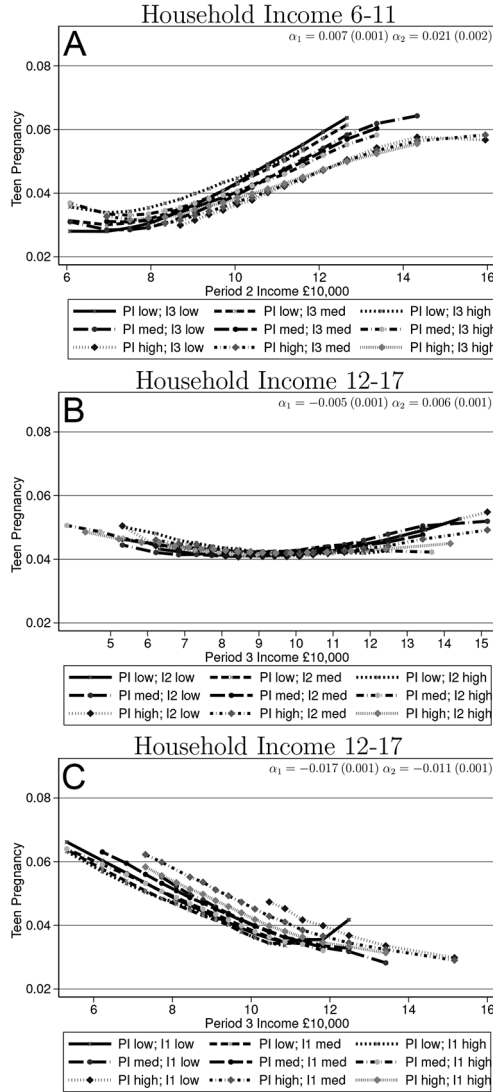


FIG. 7.—Semiparametric estimates; dependent variable is teen pregnancy; income in 2000 prices, £10,000s. Semiparametric regression controls and interpretation of α_1 and α_2 and associated standard errors are the same as in fig. 1. Low, medium, and high levels of permanent income (PI) I_1 , I_2 , and I_3 indicate income at the third, fifth and seventh deciles, respectively. Adjustments of the scale of the x-axis and, for B, the y-axis are described in sec. V.A.

cognitive ability. Teenage pregnancy, measured only for girls, may reflect instead behavior problems, excessive risk taking, or low ability to plan for the future. It is again remarkable that, with one exception, the patterns we uncover for the impact of the timing of income on these measures

of child human capital are very similar compared to those for years of schooling and earnings at age 30.

Figure 6 concerns IQ. Panels A (I_2 vs. I_1) and C (I_3 vs. I_2) are similar to those seen above for schooling and earnings, showing that I_1 and I_3 are both relatively more productive than I_2 . Panel B (I_3 vs. I_1) shows mainly increasing curves (whereas we saw inverse-U-shaped curves for schooling and earnings), suggesting that I_3 is relatively more important than I_1 . This is the only instance where there is a qualitative difference in the patterns we uncover for schooling, earnings, IQ, and teenage pregnancy.

Figure 7 concerns teenage pregnancy and shows qualitatively the same patterns we saw for schooling and earnings, although the signs of the slopes are reversed, since we see teenage pregnancy as a reflection of low human capital. The U-shaped pattern between I_1 and I_3 is quantitatively less pronounced than the inverse-U-shaped pattern for other outcomes, but it is still statistically significant. In sum, we see broadly similar results when we study the relationship between the timing of family income and very different child outcomes.

VI. Interpretation of Our Results

A. *Exogeneity of the Timing of Income Shocks*

In our main empirical specification, we cannot guarantee that the error term is orthogonal to the timing of income, even after controlling for permanent income. There are challenges to interpreting our estimates of the impact of the timing of income on the human capital of children as causal. We discuss these briefly in this section and present a much more detailed treatment of this issue (and of all the results summarized below) in appendix A3.

First, the precipitant of the income shock that we estimate, and not the income shock itself, may be the driver of the estimated effects on child outcomes. One particular example of this issue is related to labor supply. The decision of a parent to exit the labor market and spend more time with her children could lower household income and directly raise child outcomes. We show, however, that the variation in the timing of income does not come from participation decisions or hours worked. Only 0.42% and 1.46% of household and fathers' net income, respectively, is explained by labor market participation, and similarly, only 3.4% and 3.2% of the variation in net household and fathers' income, respectively, is explained by variation in hours worked (see app. A3.1).

Second, our results are robust to using exogenous variation in the timing of income generated by changes in the mapping between gross and net income, caused by reforms to the tax and benefit system between 1971 and 1998 (e.g., changes to the thresholds of tax brackets). Changes

in the tax and benefit system mean that two individuals with the same gross income profile may have very different net income profiles, depending on their cohorts. Such variation in the timing of net income, conditional on gross income, is plausibly unrelated to unobservable traits correlated with the timing of parental income and child outcomes, where identification comes from the cohort variation in the exposure to these changes (see app. A3.1).

Third, one could worry that our finding that middle-childhood income is relatively less productive than income in later adolescence is driven by the omission from our model of permanent income after the age of 17, which is highly correlated with income in adolescence. For example, the ability of parents to borrow against future income (say between the child ages of 18 and 30) is likely to be higher at age 12–17 than at age 6–11. However, there is hardly any change to our results when we carefully expand the model to include the whole life cycle of families, by taking family income either for the periods when the child is between ages 0 and 30 or for the periods when parents are between ages 25 and 55 (see app. A3.2).

Fourth, the timing of income could be a choice variable, potentially correlated with human capital investments. One obvious case is maternity leave: parents choose to take time off at the beginning of the child's life, presumably with the goal of improving the outcomes of the child (e.g., Carneiro, Løken, and Salvanes 2015). However, excluding income in the early years, when the labor supply of mothers (and to a lesser extent fathers) is affected by child-rearing activities, leads to the same conclusions as our benchmark. We also conduct two placebo tests that further reinforce our belief that, in our very rich specification, we can consider the timing of income as exogenous. We show that the timing of income does not matter for low birth weight—an outcome that predates the income. In addition, we also show that shocks to parental income occurring when the child was 30 years of age do not affect her education outcomes (see app. A3.3).

Finally, there may be unobserved traits of individuals that are potentially correlated with human capital investments and with income profiles. Those parents who have their income front-loaded in the early years of the life of the child could be different from those who have their income back-loaded. The question is, In which ways are they different?¹⁴ Our results are robust to several checks, including using different measures for the income growth profile and controlling for income volatility (see apps. A3.4 and A3.5).¹⁵

¹⁴ Parental education affects the slope of income profiles and the outcomes of children, but this is a variable we control for.

¹⁵ Our results are robust to several specification checks, which vary the set of controls, the measurement of controls, the discount rates to construct per period income, and the choice of bandwidth in the nonparametric estimation.

B. Insurance

The timing of income is important in a world where households are not able to fully insure themselves against unanticipated income shocks.¹⁶ As shown in Cunha and Heckman (2007) and Caucutt and Lochner (2012), if families face lending or borrowing restrictions, then two families with the same statistical permanent income but different time profiles of income may also have different time profiles of investments in children, resulting in children with different levels of human capital. This result, which is the basis for the empirical procedure implemented in this paper, is consistent with the standard dynamic model of parental investments in children with uncertainty.

Beyond self-insurance possibly available through access to credit markets, there are other types of insurance mechanisms available to households that we can think of as being particularly relevant to our analysis. The most obvious one is the tax and benefit system, which in Norway is quite progressive. This implies that parents are more likely to be insured against gross income shocks than against net income shocks and that the timing of parental gross income should matter less than the timing of parental net income. Blundell, Graber, and Mogstad (2015) find that the tax and benefit system provides substantial insurance against gross income shocks. Since all our previous estimates are obtained using net household incomes, we repeated our analysis, using gross incomes, and found that the resulting curves are generally flatter than the curves in figure 2, suggesting that the tax and benefit system does provide some insurance (see fig. A35). The trade-offs between I_2 and I_1 and between I_3 and I_2 are definitely flatter when we look at gross income profiles (fig. A35) than those for net income profiles (fig. 2). In fact, the relationship between child's schooling and I_2 relative to I_1 is even inverse-U shaped when we consider gross income, rather than declining when we use net income (perhaps because progressive income taxation means that a unit of gross income at a very low level of income is worth more than a unit of gross income at a high level of income). This can be seen clearly by comparing α_1 and α_2 in figure A35, which represent the difference between schooling when income moves from the bottom to the midpoint of the distribution and from the midpoint to the top of the distribution, respectively. Both α_1 and α_2 are lower when the effect of gross income is estimated. The trade-offs between I_3 and I_1 for gross income profiles are qualitatively and quantitatively more similar to those for net income profiles, relative to the corresponding trade-offs for consecutive time periods, which would suggest that the tax and transfer system does not provide too much insurance to this trade-off, perhaps because it is difficult to insure between periods that are well apart from each other, compared to periods that are close together.

¹⁶ Appendix A4 provides a full discussion of the analysis.

It is also interesting to consider the potential role of spousal insurance by comparing results using father and household incomes. If spousal insurance was important, then we would estimate larger effects of household income shocks relative to father income shocks, since the former cannot be insured through this mechanism. This is, in fact, what we find (see fig. A36, which has the same shapes but less pronounced slopes than fig. 2).

The importance of within-family insurance is also visible when we compare the role of the timing of income across siblings within the same family. Parents may be able to partly insure investments in a specific child against negative income shocks affecting that child in a particularly sensitive period, by shifting investments away from other children in the household. Indeed, when we use only within-sibling variation, we find no evidence of a relationship between the timing of income fluctuations and child outcomes (see fig. A37). This suggests that there is substantial within-family insurance against income shocks, since our estimates using cross-family variation in the timing of income show much more variation in child outcomes.

C. Simulating Large Changes to the Tax and Benefit System

We estimated a simple model of parental investments in children to simulate the extent to which changes in the progressivity of the tax and benefit system influence the way income fluctuations over the life cycle of the child affect human capital formation. With a model, one can ask which level of progressivity would be necessary to completely smooth out the effect of the timing of income shocks and how it compares with the current system and with large changes to the tax and benefit system in Norway over the past 40 years.

The structural model we consider is very simple, so we refrain from providing a detailed description here (it is described in detail in app. A5). There are three periods of childhood, parents optimally choose time and goods investments in each period, as well as employment, in a world where there is imperfect insurance against income shocks, so that investments in children react to these shocks. We estimate the technology of skill formation, assumed to be a constant elasticity of substitution production function, where the output is years of education of the child and the inputs are investments in each period. The estimated elasticity of substitution between investments in different time periods is such that investments in the first and third periods of childhood would be more productive than investments in the second period, supporting the idea of the presence of dynamic complementarities in investments (the model fits both patterns of human capital and labor market participation quite well).

Within this model we simulate how the timing of income affects human capital formation under the following alternative tax/benefit systems: (1) a system less progressive than the Norwegian system, such as that in the United States, (2) a system slightly more progressive than Norway's, and (3) a system with a very extreme degree of progressivity. We find that across a wide range of degrees of progressivity of the tax system, the timing of income shocks is always important, unless we make progressivity unrealistically large (as in case 3).

VII. Conclusion

This paper examines the importance of the timing of income shocks for the human capital development of children. Using a very large data set, consisting of the entire population of children born in Norway between 1971 and 1980, we estimate semiparametric regressions of human capital outcomes of children (measured in their adult years) on the discounted household income for the years when the child was between 0 and 17 years of age (which we label permanent income) and on income in different periods of childhood.

Using a method that controls for permanent income across the entire period of childhood in essence compares one family with a counterfactual family whose permanent income is the same, but with a different configuration of the timing of income shocks. The paper asks whether child development responds to this differential timing of family income.

We show that, keeping permanent income and income in late childhood fixed, child outcomes such as years of schooling and earnings by age 30 improve if a family experiences relatively high (counterfactual) income in early childhood and low income in middle childhood, rather than a relatively low income in early childhood and high income in middle childhood. Stating it more simply, child outcomes improve as (counterfactual) family income rises in early childhood and falls in middle childhood. We also show that a child's schooling and earnings increase when her family income falls in the middle years and increases toward the later years of childhood. Compared to the early and late period of childhood, income in the middle period seems to have low productivity.

Finally, children in families with a balanced profile of income between early and late childhood (and keeping fixed permanent income) experience better adult outcomes than those in families with income mainly "front-loaded" in the early period and those in families with income mainly "back-loaded" in the later period of childhood. In other words, we find that there is an inverse-U-shaped relationship between income in the early period, relative to income in the late period (which means keeping fixed permanent income and income in middle childhood), and the child's education and labor market outcomes in the adult years.

These empirical patterns are consistent with a model with income uncertainty, partial insurance (Blundell, Pistaferri, and Preston 2008), and complementarity between investments in children across periods (Cunha and Heckman 2007; Cunha, Heckman, and Schennach 2010). Income uncertainty, together with partial insurance possibilities and the technology of skill accumulation, leads to a setting where investments in children react to parental income shocks. Complementarity between investments taking place in different periods means that human capital is maximized when there is a balanced flow of investments. The combination of these factors can result in a model where a balanced flow of income shocks may lead to higher human capital than an unbalanced history of shocks. This would be consistent with an inverse-U-shaped relationship between the education of the child and the amount of income that is front-loaded (back-loaded) in the initial (late) period of childhood. However, our results also suggest that income in the middle period of childhood may be relatively unproductive, compared to income in earlier and later periods.

References

- Aakvik, Arild, Kjell G. Salvanes, and Kjell Vaage. 2005. "Educational Attainment and Family Background." *German Econ. Rev.* 6 (3): 377–94.
- Autor, David H., Mark Duggan, Kyle Greenberg, and David S. Lyle. 2016. "The Impact of Disability Benefits on Labor Supply: Evidence from the VA's Disability Compensation Program." *American Econ. J.: Appl. Econ.* 8 (3): 31–68.
- Baker, Michael, and Kevin Milligan. 2016. "Boy-Girl Differences in Parental Time Investments: Evidence from Three Countries." *J. Human Capital* 10 (4): 399–441.
- Barth, Erling, Karl O. Moene, and Fredrik Willumsen. 2014. "Reprint of 'The Scandinavian Model—An Interpretation'." *J. Public Econ.* 127:17–29.
- Becker, Gary S., and Nigel Tomes. 1979. "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility." *J.P.E.* 87 (6): 1153–89.
- . 1986. "Human Capital and the Rise and Fall of Families." *J. Labor Econ.* 4(3 Pt. 2): S1–S39.
- Bertrand, Marianne, and Jessica Pan. 2013. "The Trouble with Boys: Social Influences and the Gender Gap in Disruptive Behavior." *American Econ. J.: Appl. Econ.* 5 (1): 32–64.
- Björklund, Anders, and Kjell G. Salvanes. 2011. "Education and Family Background: Mechanisms and Policies." In *Handbook of the Economics of Education*, vol. 3, edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, 201–47. Amsterdam: North-Holland.
- Black, Sandra E., and Paul J. Devereux. 2011. "Recent Developments in Intergenerational Mobility." In *Handbook of Labor Economics*, vol. 4B, edited by David Card and Orley Ashenfelter, 1487–541. Amsterdam: North-Holland.
- Black, Sandra E., Paul J. Devereux, Katrine V. Løken, and Kjell G. Salvanes. 2014. "Care or Cash? The Effect of Child Care Subsidies on Student Performance." *Rev. Econ. and Statis.* 96 (5): 824–37.
- Blundell, Richard, Michael Graber, and Magne Mogstad. 2015. "Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family." *J. Public Econ.* 127:58–73.

- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *A.E.R.* 98 (5): 1887–921.
- Carneiro, Pedro, and James J. Heckman. 2003. "Human Capital Policy." In *Inequality in America: What Role for Human Capital Policies?* edited by Benjamin M. Friedman, 77–239. Cambridge, MA: MIT Press.
- Carneiro, Pedro, Katrine V. Løken, and Kjell G. Salvanes. 2015. "A Flying Start? Maternity Leave Benefits and Long-Run Outcomes of Children." *J.P.E.* 123 (2): 365–412.
- Carneiro, Pedro, Kjell Salvanes, and Emma Tominey. 2015. "Family Income Shocks and Adolescent Human Capital." Manuscript.
- Caucutt, Elizabeth M., and Lance Lochner. 2012. "Early and Late Human Capital Investments, Borrowing Constraints, and the Family." Working Paper no. 18493 (October), NBER, Cambridge, MA.
- Cunha, Flávio. 2013. "Investments in Children When Markets Are Incomplete." Manuscript.
- Cunha, Flávio, and James J. Heckman. 2007. "The Technology of Skill Formation." *A.E.R.* 97 (2): 31–47.
- Cunha, Flávio, James J. Heckman, and Susanne M. Schennach. 2010. "Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Econometrica* 78 (3): 883–931.
- Dahl, Gordon B., Katrine V. Løken, Magne Mogstad, and Kjell V. Salvanes. 2016. "What Is the Case for Paid Maternity Leave?" *Rev. Econ. and Statis.* 98 (4): 655–70.
- Duncan, Greg J., W. Jean Yeung, Jeanne Brooks-Gunn, and Judith R. Smith. 1998. "How Much Does Childhood Poverty Affect the Life Chances of Children?" *American Sociological Rev.* 63 (3): 406–23.
- Fjærli, Erik, and Rolf Aaberge. 1999. "Tax Reforms, Dividend Policy and Trends in Income Inequality: Empirical Evidence Based on Norwegian Data." Discussion Paper no. 284, Res. Dept., Statis. Norway.
- Havnes, Tarjei, and Magne Mogstad. 2011. "No Child Left Behind: Subsidized Child Care and Children's Long-Run Outcomes." *American Econ. J.: Econ. Policy* 3 (2): 97–129.
- Hilger, Nathaniel G. 2016. "Parental Job Loss and Children's Long-Term Outcomes: Evidence from 7 Million Fathers' Layoffs." *American Econ. J.: Appl. Econ.* 8 (3): 247–83.
- Humlum, Maria Knoth. 2011. "Timing of Family Income, Borrowing Constraints, and Child Achievement." *J. Population Econ.* 24 (3): 979–1004.
- Jenkins, Stephen P., and Christian Schluter. 2002. "The Effect of Family Income during Childhood on Later-Life Attainment: Evidence from Germany." Discussion Paper no. 604, Inst. Study Labor, Bonn.
- Lee, Yean-Ju, William Parish, and Robert J. Willis. 1994. "Sons, Daughters, and Intergenerational Support in Taiwan." *American J. Sociology* 99 (4): 1010–41.
- Levy, Dan, and Greg J. Duncan. 2000. "Using Sibling Samples to Assess the Effect of Childhood Family Income on Completed Schooling." Discussion paper, Northwestern Univ./Univ. Chicago Joint Center Poverty Res.
- Lindbeck, Assar. 1997. "The Swedish Experiment." *J. Econ. Literature* 35 (3): 1273–319.
- Markussen, Simen, and Knut Røed. 2014. "The Impacts of Vocational Rehabilitation." Discussion Paper no. 7892, Inst. Study Labor, Bonn.
- OECD (Organisation for Economic Co-operation and Development). 1998. *OECD Economic Outlook*, no 63. Paris: OECD.

- Pitt, Mark M., Mark R. Rosenzweig, and M. Nazmul Hassan. 1990. "Productivity, Health, and Inequality in the Intrahousehold Distribution of Food in Low-Income Countries." *A.E.R.* 80 (5): 1139–56.
- Quisumbing, Agnes. 1994. "Improving Women's Agricultural Productivity As Farmers and Workers." Discussion Paper no. 37, Educ. and Soc. Policy Dept., World Bank, Washington, DC.
- Rosenzweig, Mark R., and T. Paul Schultz. 1982. "Market Opportunities, Genetic Endowments, and Intrafamily Resource Distribution: Child Survival in Rural India." *A.E.R.* 72 (4): 803–15.
- Solon, Gary. 1999. "Intergenerational Mobility in the Labor Market." In *Handbook of Labor Economics*, vol. 3A, edited by Orley C. Ashenfelter and David Card, 1761–800. Amsterdam: North-Holland.
- Thoresen, Thor O. 2004. "Reduced Tax Progressivity in Norway in the Nineties: The Effect from Tax Changes." *Internat. Tax and Public Finance* 11:487–506.