



Intergenerational Altruism and Transfers of Time and Money: A Life-cycle Perspective

Uta Bolt, Eric French, Jamie Hentall Maccuish, and Cormac O'Dea



Intergenerational Altruism and Transfers of Time and Money: A Life-cycle Perspective

Uta Bolt

University College London and Institute for Fiscal Studies

Eric French

University College London and Institute for Fiscal Studies

Jamie Hentall Maccuish

University College London

Cormac O’Dea

Yale University and Institute for Fiscal Studies

April 2018

Michigan Retirement Research Center

University of Michigan

P.O. Box 1248

Ann Arbor, MI 48104

www.mrrc.isr.umich.edu

(734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium through the University of Michigan Retirement Research Center Award RRC08098401-09. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Regents of the University of Michigan

Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; Ron Weiser, Ann Arbor; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*

Intergenerational Altruism and Transfers of Time and Money: A Life-cycle Perspective

Abstract

Parental investments in children can take one of three broad forms: (1) Time investments during childhood and adolescence that aid child development, and in particular cognitive ability; (2) educational investments that improve school quality and hence educational outcomes; (3) cash investments in the form of inter vivos transfers and bequests. We develop a dynastic model of household decision-making with intergenerational altruism that nests a child production function, incorporates all three of these types of investments, and allows us to quantify their relative importance and estimate the strength of intergenerational altruism. Using British cohort data that follows individuals from birth to retirement, we find that around 40% of differences in average lifetime income by paternal education are explained by ability at age 7, around 40% by subsequent divergence in ability and different educational outcomes, and around 20% by inter vivos transfers and bequests received so far.

Citation

Bolt, Uta, Eric French, Jamie Hentall Maccuish, and Cormac O’Dea. 2018. “Intergenerational Altruism and Transfers of Time and Money: A Life-cycle Perspective.” Ann Arbor MI: University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2018-379. <http://mrrc.isr.umich.edu/wp379/>

Authors’ acknowledgements

Preliminary, please do not quote. For excellent research assistance we thank Jack Light and for helpful comments we thank Andrew Hood, George Levi-Gayle, Limor Golan, and Ananth Seshadri. Funding from the Social Security Administration through the Michigan Retirement Research Center (MRRC grant UM17-15) and the Economic and Social Research Council (Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (RES-544-28-50001), and Grant Inequality and the insurance value of transfers across the life cycle (ES/P001831/1)) for this work is gratefully acknowledged. O’Dea additionally acknowledges funding from the ESRC Secondary Data Analysis grant ref: ES/N011872/. Correspondence to eric.french.econ@gmail.com or cormac.odea@yale.edu. Any errors are the authors’ own.

1 Introduction

Intergenerational links are a key determinant of levels of inequality and social mobility, with previous work looking at a range of developed economies finding very significant intergenerational correlations in education, incomes and wealth (e.g. Dearden et al. 1997; Mazumder 2005; Charles and Hurst 2003; Chetty et al. 2014). The literature on understanding the mechanisms behind this persistence is much newer. Understanding the drivers of this persistence of economic outcomes across generations is crucial for the design of tax and transfer policies for two main reasons. First, insofar as the correlations reflect differential parental investments in children (both of time and money) they represent an important reason that the design of public policy should not treat the distributions of ability, education, earnings, and wealth as fixed. Policies designed to mitigate the intergenerational transmission of inequality through one channel (e.g., estate taxes) could, by affecting parental investments, increase transmission through another channel (e.g., parental spending on children's education). Second, the extent of parental investment in children over the course of their lives provides important evidence on the extent of intergenerational altruism — the extent to which parents forgo consumption and leisure to invest in their children allows us to estimate the relative weight they put on their children's welfare relative to their own. The degree of intergenerational altruism is a key parameter for assessing the potential benefits of social security and tax reform, since current generations will only be willing to accept cuts to their benefits in order to reduce budget deficits if they are altruistic towards future generations (Fuster et al. 2007).

In this paper, we develop a dynastic model of household decision-making that incorporates three different types of parental investment in children: i) time investments during childhood and adolescence that aid child development, and in particular cognitive ability, ii) educational

investments that improve school quality and hence educational outcomes, and iii) cash investments in the form of inter vivos transfers and bequests. The key contribution of the paper is to estimate such a model using unique longitudinal data from a survey that has been running for 60 years — following a cohort of individuals from birth to retirement. Using these data, we can measure parental inputs over the whole life cycle, and hence look directly at early life investments of time and goods, estimate a child production function for cognitive ability and link that ability measure to earnings in adulthood. The data also include detailed information about the schooling received by individuals and the inter vivos transfers they then receive from parents during early adult life.

Using these data, we are able to build and estimate a model capable of speaking to the issues raised above. First, we can provide an estimate of the degree of intergenerational altruism drawing on data on a number of different investment decisions. Such an estimate is likely to be more robust than one based on a single decision (such as how much to leave in bequests) which is likely to be affected by a number of other confounding factors. Second, having estimated the degree of intergenerational altruism (along with the other structural parameters that govern household behaviour), we can run policy counterfactuals and look at how each type of parental investment would respond.

Preliminary analysis of this cohort data suggests that around 40% of differences in average lifetime income by paternal education are explained by ability at age 7, around 40% by subsequent divergence in ability and different educational outcomes, and around 20% by inter vivos transfers and bequests received so far. These findings are supported by results from a simple version of the model that has been calibrated to match wealth and labour supply moments. Using consumption equivalent variation to measure the welfare gains from higher-

educated parents, we again find that differences in investments before and after the age of 7 are of roughly equal importance in determining lifetime utility differences between children of high- versus low-educated parents, with investments in ability and education looking much more important than differences in the level of inter vivos transfers and bequests. Looking in more detail at investments in ability, we find that higher levels of time investments increase ability, and that the ability production function looks to exhibit dynamic complementarity, at least at younger ages (see Cunha et al. (2010)).

Finally, we present estimates of many of the investments that households make in their children, including time and money investments. We show that increased investment of time and goods of parents leads to higher ability children (as measured by test scores), and this higher ability leads to higher wages and incomes later in life. Furthermore, we show that higher income parents invest more in their children, and that these investments can explain much of the difference in lifetime incomes of children across the parental education distribution.

This paper relates to a number of different strands of the existing literature, including work measuring the drivers of inequality and intergenerational correlations in economic outcomes, the large literature seeking to understand child production functions and work on parental altruism and bequest motives. The most closely related papers, however, are those focused on the costs of and returns to parental investments in children. Our paper is most similar to Lee and Seshadri (2016). They develop a model that has many similar features to that used in our paper, but they lack data that links investments at young ages to earnings at older ages. As a result, they have to calibrate key parts of the model, while we are able to estimate the human capital production technology, and show how early life investments and the resulting human capital impacts late life earnings. Caucutt and Lochner (2012) is also related to ours. Their paper estimates a human

capital production function and altruistic parental transfers to improve human capital of children. They find that borrowing constraints are an important deterrent to college going. They use data on parental investments at different ages and also later life income, but, unlike us, they cannot directly measure early life ability. Furthermore, they restrict the set of investments that can be made in children because they do not allow for endogenous labor supply or inter vivos cash transfers. Other closely related papers include Del Boca et al. (2014) and Gayle et al. (2015), both of which develop models in which parents choose how much time to allocate to the labour market, leisure, and investment in children. Neither paper, however, incorporates household savings decisions, and hence, the tradeoff between time investments in children now and cash investments later in life. Abbott et al. (2016) focuses on the interaction between parental investments, state subsidies, and education decisions, but abstract away from the role of parents in influencing ability prior to the age of 16. Castaneda et al. (2003) and De Nardi (2004) build overlapping-generations models of wealth inequality that includes both intergenerational correlation in human capital and bequests, but does not attempt to model the processes underpinning the correlation in earnings across generations.

The rest of this paper proceeds as follows. Section 2 describes the data, and documents descriptive statistics on ability, education, and parental investments. Section 3 lays out the dynastic model used in the paper. Section 4 then provides some reduced-form evidence on the impact of parental investments, before Section 5 provides some initial results on the relative importance of different channels in explaining intergenerational correlations in education, earnings and welfare. Section 6 concludes, and draws out some implications for policy.

2 Data and descriptive statistics

The key data source for this paper is the National Child Development Study (NCDS). The NCDS follows the lives of all people born in England, Scotland, and Wales in one particular week of March 1958. The initial survey at birth has been followed by subsequent follow-up surveys at the ages of 7, 11, 16, 23, 33, 42, 46, 50, and 55.¹ During childhood, the data includes information on a number of ability measures, measures of parental time investments (discussed in more detail below) and parental income. Later waves of the study record educational outcomes, receipt of inter vivos transfers, demographic characteristics, earnings, and hours of work. For the descriptive analysis in this section, we focus on those individuals for whom we observe both their father's educational attainment (age left school) and their own educational qualifications by the age of 33. This leaves us with a sample of 9,436 individuals.

The main limitation of the NCDS data currently available for our purposes is that we do not have data on the inheritances received or expected by members of the cohort of interest. We therefore supplement the NCDS data using the English Longitudinal Study of Ageing (ELSA). This is a biennial survey of a representative sample of the 50-plus population in England, similar in form and purpose to the Health and Retirement Study (HRS) in the U.S. The 2012-13 wave of ELSA recorded lifetime histories of inheritance receipt, and since we also observe father's education in those data, we can use those recorded receipts to augment our description of the divergence in lifetime economic outcomes by parental background. We focus on individuals in ELSA born in the 1950s, leaving us with a sample of 3,001.²

¹ The age-46 survey is not used in any of the subsequent analysis as it was a telephone interview only, and the data are known to be of lower quality.

² The next wave of the NCDS, which will be in the field next year, is currently planned to collect information on lifetime inheritance receipt. We hope to use these new data in later versions of this work.

In the rest of this section, we document the evolution of inequalities over the life cycle, and in particular how they relate to parental background and parental investments over time.

2.1 Ability and time investments

We have reading and math test scores for our cohort of interest at the ages of 7, 11 and 16. At each age we create our preferred measure of individual ability by taking the average of the percentage score on each test, and then normalise to ease interpretation.


Figure 1 shows the cumulative distribution of normalised ability at each age, splitting the sample according to father's education (compulsory only, some post-compulsory, some college - the proportion of children in each group is shown in the Appendix Table). For this age group of fathers, compulsory education roughly corresponds to leaving school at age 14, post-compulsory means leaving school between ages 15 and 18, and some college means staying at school until at least age 19. It shows that, as one might expect, children whose fathers have a higher level of education have higher ability; at the age of 7, 23% of the children of low-education fathers had ability around one standard deviation or more below the mean, compared to just 2% of the children of high-education fathers. Similarly, 22% of the children of high-education fathers had ability around one standard deviation or more above the mean, compared to 7% of the children of low-education fathers.

The second key thing to note from Figure 1 is that ability gaps by father's education widen through childhood. At the age of 7, 44% of the children of low-education fathers have above-average ability³ compared to 76% of the children of high-education parents - a gap of 32 percentage points. By age 11, that gap has widened to 42 percentage points, and by age 16 it stands at 46 percentage points.

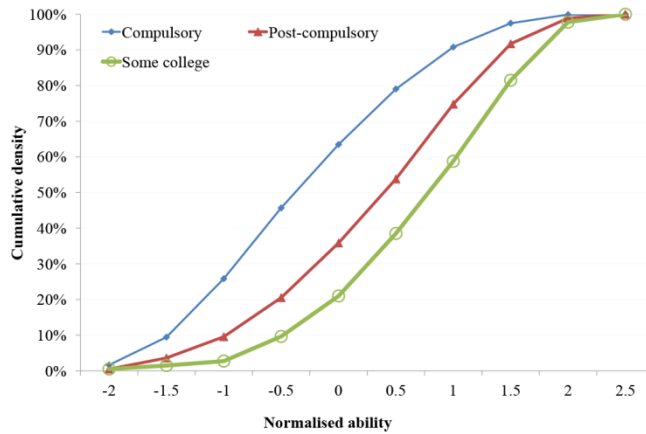
³ That, is their normalised ability is greater than 0.25.

Figure 1: Normalised ability at age 7, by parental education


(a) Age 7



(b) Age 11



(c) Age 16



Tables 1, 2, and 3 provide some descriptive evidence that at least some of the widening in ability gaps by parental characteristics between ages 7 and 16 age can be explained by differential parental investments (we investigate this hypothesis more formally in Section 4). Table 1 documents parental responses to a question about reading with their child, asked when the child is 7. It shows relatively small, but potentially important, differences in the frequency with which both mothers and fathers read to their children, splitting families according to the education of the father. For example, 34% of fathers with only compulsory education read to their 7-year-old children each week, compared to 53% of fathers with some college education.

Tables 2 and 3 present the child's teachers assessment of parental interest in the child's education, at the ages of 7 and 11 respectively. The differences by father's educational attainment are perhaps even more striking than those in reading patterns. When the child is 7, fathers with some college education are three times more likely to be judged by the teacher to be 'very interested' in their child's education as fathers with just compulsory education (65% compared to 22%). At the age of 11, the gap in paternal interest is very similar, with 72% of college-educated fathers judged to be 'very interested' in their child's education, compared to 25% of fathers with just compulsory education. The tables also show that having a higher-educated father dramatically reduces the risk of a child having parents with little interest in their education. Among those with a college-educated father, only around 10% have a mother or father who is judged to show 'little interest' in their education at the age of 11. On the other hand, among those whose father has only compulsory education that figure rises to around a quarter of mothers and nearly half of fathers.

Table 1: Frequency with which parents read to age-7 children

		Father reads...		
		Never	Sometimes	Every week
Father's education	Compulsory	30%	36%	34%
	Post-compulsory	20%	35%	45%
	Some college	18%	29%	53%
		Mother reads...		
		Never	Sometimes	Every week
Father's education	Compulsory	16%	37%	47%
	Post-compulsory	12%	31%	57%
	Some college	10%	23%	67%

Table 2: Teacher assessment of parental interest in education of age-7 child

		Father		
		Very interested	Some interest	Little interest
Father's education	Compulsory	22%	24%	55%
	Post-compulsory	44%	22%	34%
	Some college	65%	15%	20%
		Mother		
		Very interested	Some interest	Little interest
Father's education	Compulsory	35%	43%	23%
	Post-compulsory	60%	30%	10%
	Some college	76%	18%	6%

Table 3: Teacher assessment of parental interest in education of age-11 child

		Father		
		Very interested	Some interest	Little interest
Father's education	Compulsory	25%	29%	46%
	Post-compulsory	54%	25%	21%
	Some college	72%	16%	12%
		Mother		
		Very interested	Some interest	Little interest
Father's education	Compulsory	35%	38%	26%
	Post-compulsory	61%	27%	12%
	Some college	76%	16%	8%

2.2 Educational attainment and school type

Table 4 shows the correlation in educational attainment between fathers and their children. It shows two dramatic impacts of paternal education on educational outcomes. First, having a high-educated father makes it much less likely that a child will end up dropping out of high school.⁴ 30% of the children of fathers with just compulsory education end up as high-school dropouts, compared to only 10% of those whose fathers have some post-compulsory education, and just 2% of those whose father have some college education. Second, having a high-educated father makes it much more likely that a child will end up with some college education. Fully 66% of the children of college-educated fathers also end up with some college education, compared to only 20% of those whose fathers only have compulsory education.

Table 4: Intergenerational correlation in education

		Child's education		
		High-school dropout	High-school graduate	Some college
Father's education	Compulsory	30%	50%	20%
	Post-compulsory	10%	47%	43%
	Some college	2%	32%	66%

Of course, it is in theory possible that all of the intergenerational correlation in education is explained by the relationship between parental education and ability documented in the previous subsection. However, one might also think that differences in the quality of the schools attended by children from different backgrounds also plays a role. In our particular institutional context (children attending high school in Britain in the late 1960s and early 1970s), a key dimension in which schools differed in quality was their 'type.' The majority of children attended

⁴ In the U.K. context, we define 'high school dropout' as not having any of the academic qualifications obtained at age 16 (formerly O-Levels, now GCSEs).

‘comprehensive’ or ‘secondary modern’ public schools that drew their students from across society (henceforth we refer to this type of school simply as comprehensive). A small proportion attended ‘grammar’ schools: public schools to which admittance was by an ability test at the age of 11. In addition to the peer effects associated with attendance at such a school, these grammar schools attracted much better teachers on average, and were much more focused on university (college) attendance than other public schools. Finally, a small minority of children went to private schools.

Table 5 shows the distribution of children across these three different types of school. As one might expect, those with higher educated fathers are dramatically more likely to have attended higher quality schools. 30% of those whose fathers went to college attended a private high school compared to just 2% of those with low-educated fathers, and a further 26% attended a grammar school, compared to just 9% of those with low-educated fathers. Of course some of this discrepancy (particularly in the case of grammar schools) might well be accounted for by the differences in ability documented above, but they also reflect differential financial investments in children’s education. The most obvious form of educational investment is paying for private education, which is much higher quality on average than public education. However, educational investments could also take less direct forms, such as paying the house price premium associated with living in the neighborhood of a good public school.

Table 5: High-school type by father’s education

		Child’s school type		
		Comprehensive	Grammar	Private
Father’s education	Compulsory	89%	9%	2%
	Post-compulsory	68%	18%	13%
	Some college	43%	26%	30%

These financial investments differ from inter vivos transfers and bequests in terms of timing, but also more importantly in that they directly impact on children's earnings through the returns to education. The relationship between school type and educational attainment, conditional on ability, is shown by Figures 2 and 3. We divide our sample into quintiles of ability (measured at age 16 in the way described above), and then plot the probability of completing high school and attending college respectively separately for individuals that attended each of the three school types.

Figure 2 shows that at all levels of ability outside the top quintile, children who attend a grammar or private school are much more likely to complete high school than those of the same ability attending a normal public school. For example, 80% of children in the middle quintile of age-16 ability at a comprehensive school complete high school, compared to around 95% of those of the same ability who attend either a grammar or private school.

Figure 3 shows that attendance at private school provides a clear boost to the probability of college attendance conditional on age-16 ability. While individuals in the middle quintile of ability who attend a comprehensive school have less than a 30% chance of ending up with some college education, those with the same ability who attended private school have more than a 40% chance.

Figure 2: The impact of school type on completing high school

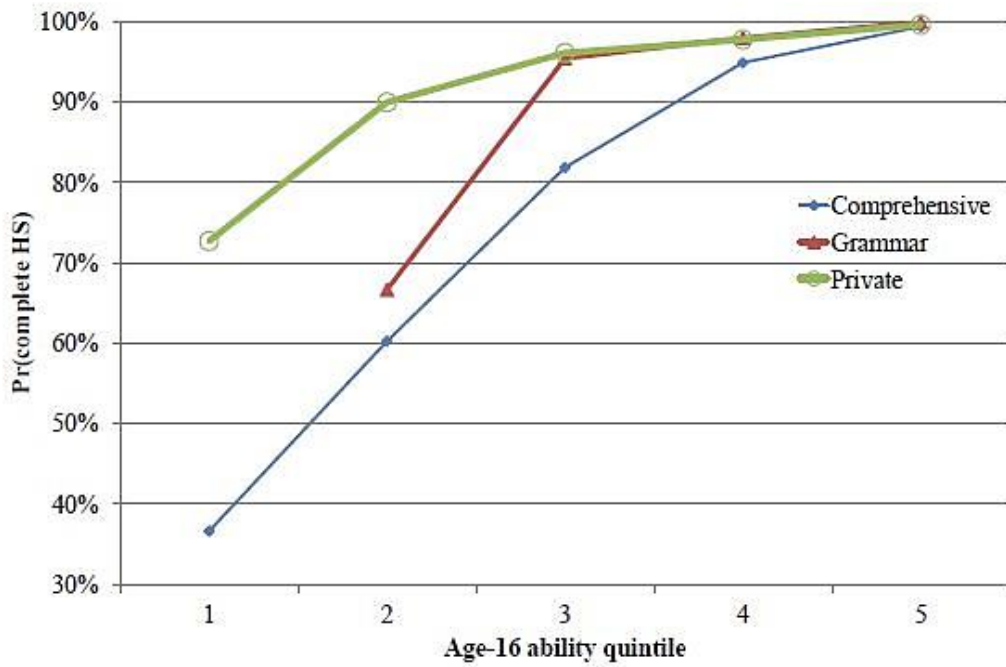


Figure 3: The impact of school type on attending college

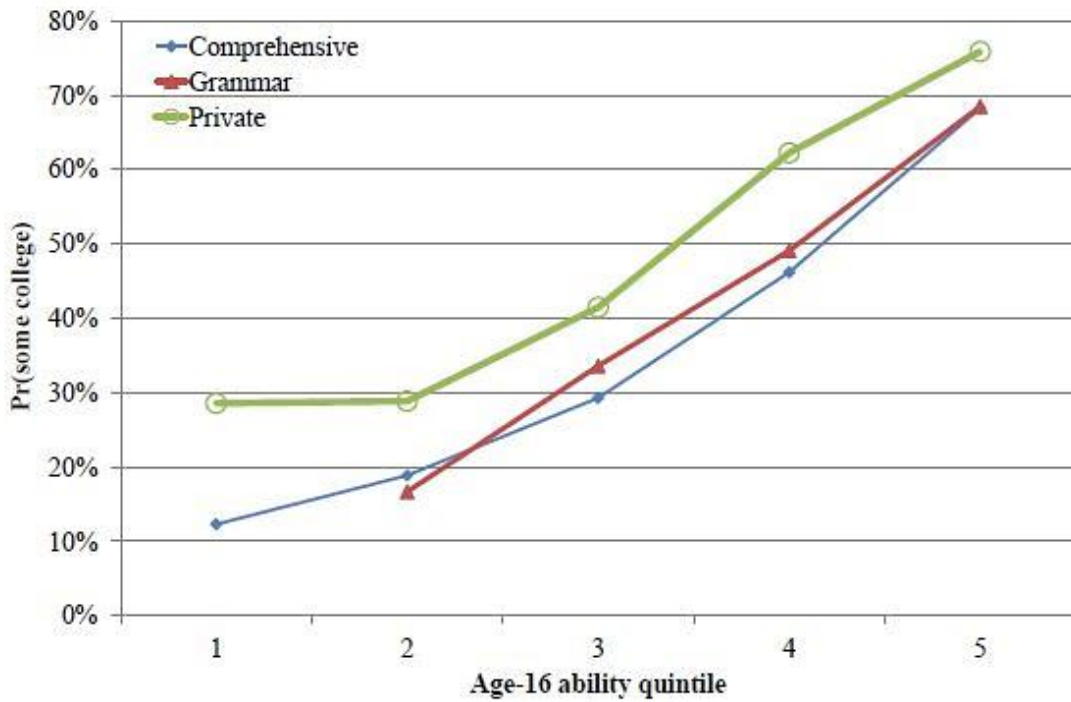


Table 6: Receipt of inter vivos transfers and bequests by father's education

		Inter vivos transfers by age 33		
		Mean (£)	Received	Mean exc. zeros (£)
Father's education	Compulsory	5,805	24%	24,281
	Post-compulsory	11,071	41%	27,008
	Some college	31,547	55%	56,933

		Inheritances (1950s birth-cohort)		
		Mean (£)	Received	Mean exc. zeros (£)
Father's education	Compulsory	17,180	26%	66,545
	Post-compulsory	43,901	40%	110,024
	Some college	55,669	46%	120,843

2.3 Inter vivos transfers and bequests

Table 6 documents the receipt of inter vivos transfers and bequests of the NCDS cohort so far, again splitting by father's education. As explained at the start of this section, the top panel draws on the NCDS data itself, while the bottom panel uses ELSA data instead, as information on inheritance receipt is not yet available in the NCDS.

The table shows that inter vivos transfers are a significant source of economic resources for young adults, and that as one would expect are much more significant for those with higher-educated parents. By the age of 33, 55% of those whose fathers attended college had received an inter vivos transfer, of an average of around £50,000. While this is the mean of a highly right-skewed distribution, these figures indicate an important role for inter vivos transfers relieving borrowing constraints in this part of the life cycle. At the same age, 24% of those with low-educated fathers had received an inter vivos transfer, of an average size of just less than £25,000.

Evidence from ELSA data suggests that differences in inheritance receipt by parental background are also significant. 46% of those with college-educated fathers have received an inheritance, compared to 26% of those with low-educated fathers, and among those who have received an inheritance, those with college-educated fathers have received around twice as much

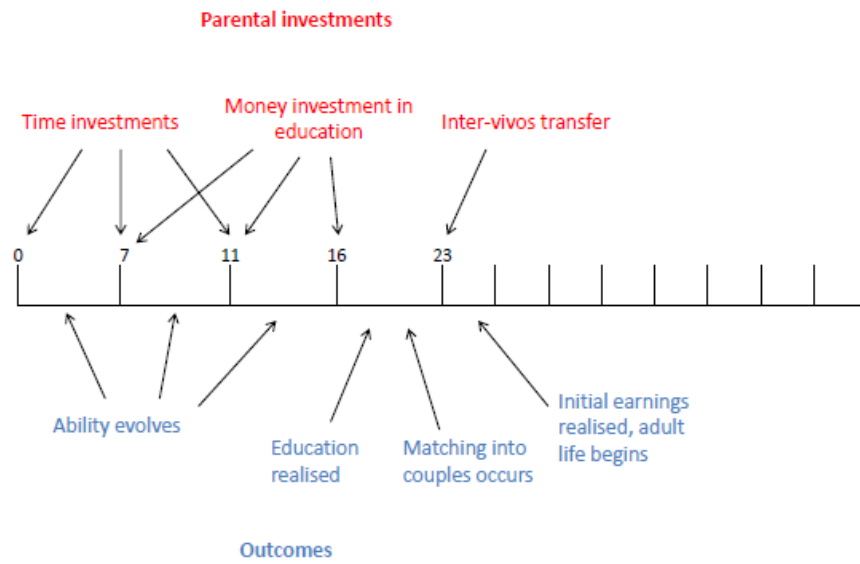
on average (£120,843 compared to £66,545) . The net result is that those with college-educated fathers have inherited around £40,000 more than those with low-educated fathers. This is likely to understate the true difference in mean lifetime inheritance receipt between these groups; some of those born in the 1950s will still have living parents, and differential mortality means it is in fact likely that this applies to a larger share of those with high-educated fathers.

3 Model

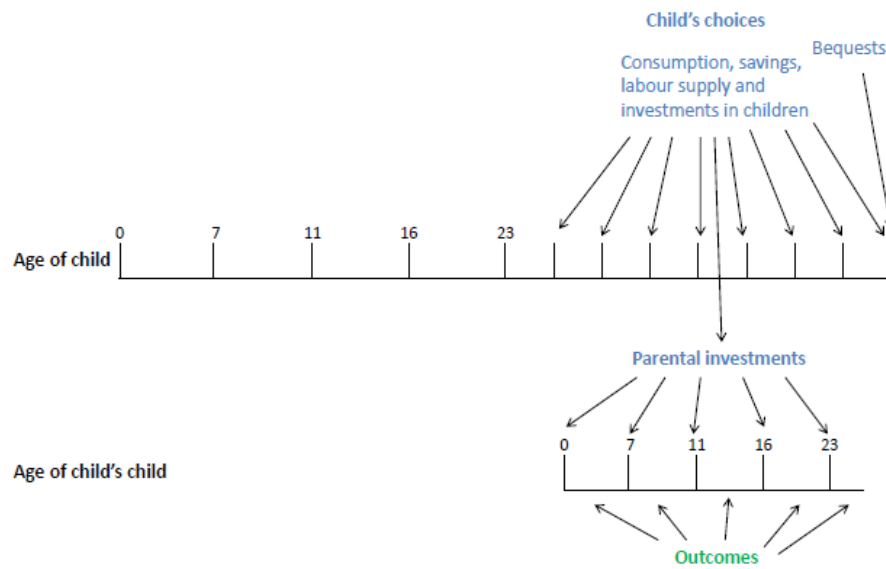
This section describes a dynastic model of consumption and labor supply in which parents can make different types of transfers to their children. The model can be used to a) evaluate how particular intergenerational transfers affect household behavior, b) compare the relative insurance value of these types of transfers, and c) simulate household behavior and welfare under counterfactual policies (for example, under reforms to estate taxation). Figure 4 provides an overview of the dynastic model. During childhood, parental time investments in children and money investments in education affect the evolution of the child's ability and their educational attainment. Children are then matched in couples, receive any inter vivos transfer from their parents and begin adult life. They then have their own children, and alongside the standard choices of consumption and labour supply they choose how much to invest in their own children, with implications for their children's future outcomes.

Figure 4: The life cycle of an individual

(a) Childhood



(b) Adulthood and Investment in Children



We now provide formal details of the model. First, we outline a production function for ability, schooling and education in Section 3.1. We then outline the decision problem of a couple with a dependent child in Section 3.2.

3.1 *A production function for ability, schooling and education*

This section describes the production function for ability, schooling, and education from ages birth to age 23. Over this part of the life cycle, the child makes no decisions. However, their parents do make decisions about the investments of time and goods received by their children. These choices do not directly impact the contemporaneous utility of the child, but leads to higher wages, incomes, and higher quality spouses later in life, which does increase their later life outcomes.

3.1.1 *Child ability production function*

A child's ability at birth is given by:

$$ab_0 = f_{ab_0}(ed^m, ed^f, u^{ab_0}) \quad (1)$$

where ed^m and ed^f represent the education levels of the child's parents (m and f index male and female respectively) and u^{ab_0} is a stochastic variable that generates heterogeneity in initial ability, conditional on parental education. Between birth and age 16, child ability updates each period according to the transition equation given in (2). The rate of growth of a child's ability depends on his/her parents' level of education, the time investments (ti_t^m, ti_t^f) those parents make, and the child's school type (st). There is also a stochastic component to the ability transition equation (u_{t+1}^{ab}).

$$ab_{t+1} = f_{ab}(ab_t, ed^m, ed^f, ti_t^m, ti_t^f, st_t, u_{t+1}^{ab}) \quad (2)$$

Ability evolves until the age of 16, after which it does not change (ab_{16} without a subscript denotes final ability).

3.1.2 School type production function

School type (st) is assumed not to vary between the start of education and the age of 11. At the age of 11, school type is realised as one of three outcomes: 1) Private ($st = p$), 2) Public – high quality ($st = g$), 3) Public – low quality ($st = m$).⁵

Parents can make one or both of two types of money investments in their children’s schooling. First, they can pay a quantity of their choosing (mi^g) to attempt to get their children into a high quality public school (one can think of this as paying a premium to locate in a district where access to good quality schools is easier). Second they can make money investments in private schooling, paying a cost ($mi^p = p$), to guarantee that their child gets into a private school. We model the outcome of the school type as following a two-stage process. First, the child’s type of public school is realised (st^g is a binary indicator of getting an *offer* at a high quality — or ‘grammar’ school). This is a stochastic function of the child’s ability, parents’ education and parent’s choosing to spend money living (mi^g) in a location where access to good schools is easier:

$$st^g = f_{st^g}(ab_{11}, ed^m, ed^f, mi^g, u^{st^g}) \quad (3)$$

⁵ This component of the model is motivated by the institutional structure that faced the cohort represented by our main data. For this cohort, children took an exam at the age of 11 (the ‘eleven-plus’). Children who performed well in this exam got a place in a selective ‘grammar school.’ Children who performed less well got a place in a ‘secondary modern’ or ‘comprehensive’ school. In our counterfactual analysis, we will explore scenarios in which there is no link between ability and the quality of public schooling.

Second, after observing the type of public schooling on offer for their child, parents decide whether or not to pay for private schooling. They can accept the public option that their child has been given and pay $mi^p = 0$ or to reject it and pay $mi^p = p$ for private schooling.

This process can be summarised as follows:

$$st = \begin{cases} m & \text{if } mi^p = 0 \text{ and } st^g = 0 \\ g & \text{if } mi^p = 0 \text{ and } st^g = 1 \\ p & \text{if } mi^p = p \end{cases}$$

Total money investments $mi = mi^g + mi^p$: the sum of payments aim at gaining access to good public schools and those aimed at securing access to private schools.

3.1.3 Education Production Function

Education takes one of three values: High School drop out, High School graduate, and Some College. It is realised in the period prior to a child turning 23. The education production function depends on the (now grown-up) child's ability, their school quality and a stochastic variable (u^{ed}).

$$ed = f_{ed}(ab, st, u^{ed})$$

3.2 Parents' decision problem

3.2.1 Stages of life

An individual's adult life cycle starts at the age of 23 at which point their ability has been formed through their parents' decisions, their education has been realised and they have been matched into couples. Their life cycle has three stages. First, there is the **early adult phase**, from the age of 23 to 48 when couples make decisions as a collective unit and have a dependent child. There is then a one-period **transition phase** at the age of 49 which is the last age at which they

make decisions on behalf of their child. From the age of 50, their child has grown up and they enter their **late adult phase**. During this phase couples are subject to stochastic mortality risk.

In outlining the dynastic model we describe below a life-cycle decision problem of a single generation. All generations are, of course, linked; each member of the couple whose decision problem we specify has parents, and they, in turn, will have children. We will refer to the generation whose problem we outline as generation 1, their parents as generation 0, and their children as generation 2. In the exposition below, model periods are indexed by the age of the members of the couples in generation 1.⁶

3.2.2 *Initial conditions and marital matching*

Individuals start the decision-making phase of their life in couples at the age of 23. Individuals differ at the start of life in their ability, their level of education and their initial wealth. The first two are generated according to the production functions with inputs determined endogenously by their parents. The third — initial wealth — will come as a cash gift from their parents (the parents' decision problem is outlined below).

Before they make any decisions, individuals are matched into couples and acquire a dependent child at the age of 26. There is probabilistic matching between men and women based on education and ability. The probability that a man of education ed^m and ability ab^m gets married to a woman with education ed^f and ability ab^f is given by $Q^m(ed^m, ab^m, ed^f, ab^f)$. The (symmetric) matching probabilities for females are $Q^f(ed^f, ab^f, ed^m, ab^m)$. Everyone is matched into couples – there are no singles in the model.

⁶ That is, subscripts are an index of calendar time, not of age. For example, $V_{50}^1()$ is the value function of generation 1 at the age of 50, but $V_{50}^2()$ is the value function of generation 2 in the year that generation 1 was 50 years old.

3.2.3 Utility and demographics

The utility of each member of the couple $g \in \{m, f\}$ (male and female respectively)

depends on their consumption and leisure:

$$u_g(c, l) = \frac{(c^{\nu_g} l^{(1-\nu_g)})^{1-\gamma}}{1-\gamma}$$

We allow the relative preferences for consumption and leisure to vary with gender.

Household preferences are given by the equally-weighted sum of male and female utility:

$$u(c, l^m, l^f) = u_m(c^m, l^m) + u_f(c^f, l^f)$$

and the consumption outcome is efficient within the household.

Mortality is stochastic — the probability of survival of a couple (we assume that both members of a couple die in the same year) to age $t + 1$ conditional on survival to age t is given by s_{t+1} . We assume that death is not possible until the household enters the late adult phase of life at the age of 50 and that death occurs by the age of 110 at the latest.

3.2.4 Discounting and intergenerational altruism

In discounting their future utility, each generation applies a discount factor (β). Each generation is altruistic regarding the utility of their offspring (and indeed future generations). In addition to the time discounting of their children's future utility (which they discount at the same geometric rate at which they discount their own future utility), they additionally discount it with an intergenerational altruism parameter (λ).

3.2.5 Decision problem in early adult phase of life

Decisions. During this phase, couples in generation 1, matched into couples are making decisions on their own behalf and on behalf of their dependent child (generation 2). They make

up to four choices each period. These are (with the time periods in which those decisions are taken given in parentheses):

1. Household consumption – c (each period)
2. Hours of work for each parent — h^m, h^f where m and f index hours of work by the male and female respectively (each period). We allow each parent to work full-time, part-time or not at all.
3. Time investments in children – ti (up to and including the age at which their child turns 11)
4. Private schooling choice (equivalently money investments in children's education) — mi (only at the age that their child turns 11)

Constraints. Parents face two types of constraints. The first is an intertemporal budget constraint at the household level

$$a_{t+1} = (1 + r)(a_t + y_t - c_t - mi_t) \quad (4)$$

where a is parental wealth, y is household income and the other variables have been defined above. Wealth must be non-negative in all periods. The second constraint is a per-parent ($g \in \{m, f\}$) intratemporal time budget constraint:

$$T = l_t^g + ti_t^g + h_t^g \quad (5)$$

where T is a time endowment, l^g is leisure time and the other variables have been defined above.

Earnings and Household income. Household income is given by $y = \tau(e^m, e^f)$, where $\tau(\cdot)$ is a function which returns net-of-tax income and e^m and e^f are male and female earnings respectively. Earnings are equal to hours multiplied by the wage rate, e.g.: $e^f = h_t^f w_t^f$. That

wage rate evolves according to a process that has a deterministic component which varies with age and a stochastic (AR(1)) component.

$$\begin{aligned}\ln w_t &= \delta_0 + \delta_1 t + \delta_2 t^2 + \delta_3 t^3 + \delta_4 \ln ab_{16} + \delta_5 PT + v_t \\ v_t &= \rho v_{t-1} + \eta_t \\ \eta &\sim N(0, \sigma^2)\end{aligned}$$

where PT is a dummy for working part time. While the associated subscripts are suppressed here, each of $\{\delta_0, \delta_1, \delta_2, \delta_3, \delta_4, \delta_5, \rho, \sigma^2\}$ varies by gender (g) and education (Ed).

Uncertainty. In this phase, couples face uncertainty over the innovation to their wage equation, over the stochastic innovations to the child ability production function and the school type production function. The joint distribution of these stochastic variables

($\mathbf{q}_t^e \equiv \{\eta_t^m, \eta_t^f, u_t^{ab}, u_t^{st}\}$) is given by $F_t^e(\mathbf{q}_t^e)$.

State Variables. The vector of state variables for generation 1 during the early adult phase of life contains (collected in the vector $\mathbf{X}^{1,e}$):

1. Age (t),
2. Assets (a^1),
3. Wage rates ($w^{m,1}, w^{f,1}$),
4. Education levels ($ed^{m,1}, ed^{f,1}$),
5. Own abilities ($ab^{m,1}, ab^{f,1}$),
6. Child's gender (g^2),
7. Child's ability (ab^2),
8. Child's school type (st^2),

where we make explicit the generation to which the state variable corresponds.

Value function. The value function for generation 1 in the early adult phase of life ($V^{1,e}$) is given below in expression (6):

$$V_t^{1,e}(\mathbf{X}_t^{1,e}) = \max_{c_t, h_t^m, h_t^f, t_i, m_i} \left(u(c_t, l_t^m, l_t^f) + \beta \int V_{t+1}^{1,e}(\mathbf{X}_{t+1}^{1,e}) dF_{t+1}^e(\mathbf{q}_{t+1}^e) \right) \quad (6)$$

s.t. i) the intertemporal budget constraint in equation (4)

ii) and the time budget constraints in equation (5)

3.2.6 Decision problem in the transition phase

The final period in which a couple is making decisions on behalf of their dependent child is when they are 46 (and their child is 23).

Decisions. During this phase, couples make three sets of choices:

1. Household consumption – c (each period);
2. Hours of work for each parent – h^m, h^f where m and f index hours of work by the male and female respectively (each period);
3. A cash gift (x) to their children. This gift represents inter vivos transfers and inheritances.

Constraints. Parents once again face two types of constraints – an intratemporal time constraint and an intertemporal budget constraint. The former is the same as that given in equation 5 in describing the early adult phase of the life cycle (except that time investments in children will now always be zero). The intertemporal budget constraint in this phase takes account of the cash gifts and is given in equation 7.

$$a_{t+1} = (1 + r)(a_t + y_t - c_t - x_t) \quad (7)$$

State variables The set of state variables ($\mathbf{X}^{1,tr}$) in this phase is that same as in the early phase of adulthood ($\mathbf{X}^{1,e}$).

Uncertainty. Couples now face two distinct types of uncertainty. The first is uncertainty regarding their own circumstances next year – that is their next period wage draws

($\mathbf{q}_{t+1}^{tr} \equiv \{\eta_{t+1}^m, \eta_{t+1}^f\}$ with distribution given by $F_{t+1}^{tr}(\mathbf{q}_{t+1}^e)$). The second is uncertainty over the characteristics of their child the following period. The dimensions of uncertainty here are the child’s education, their initial wage draw, and the attributes of their future spouse (his/her ability, education level, assets, and initial wage draw). The stochastic variables are collected in a vector \mathbf{p}_{t+1} , and their joint distribution is given by $H()$.

Value function. The decision problem of generation 1 in the transition phase of life is:

$$V_t^{1,tr}(\mathbf{X}_t^{1,tr}) = \max_{c_t, h^m, h^f, x} \left(u(c_t, l_t^m, l_t^f) + \beta \int V_{t+1}^{1,l}(\mathbf{X}_{t+1}^{1,l}) dF_{t+1}(\mathbf{q}_{t+1}) + \beta\lambda \int V_{t+1}^{2,e}(\mathbf{X}_{t+1}^{2,e}) dH(\mathbf{p}_{t+1}) \right) \quad (8)$$

s.t. the intertemporal budget constraint in equation (7)

and the time budget constraint in equation (5)

Note that there are two continuation value functions here. The first is the future expected utility of that the decision-making couple will enjoy in the next period (when they will enter the late adult phase). The value function (given in equation (9)) must be integrated with respect to next period’s wage draws, which are stochastic, and discounted by β , the time discount factor. The second continuation value function is the expected value of the couple to which the child of the generation 1 decision-maker will belong. The (altruistic) parents take this into account in making their decisions. This continuation utility is discounted by both the time discount factor

and the altruism parameter (λ). This value function is the early adult value function for generation 2 (the equivalent for generation 2 of the value function given in equation (6)).⁷

3.2.7 Decision problem in the late adult phase

At this stage the children of generation 1 have entered their own early adult phase and the generation 1 couple enters a ‘late adult phase’,

Decisions. During this phase households make labor supply and consumption/saving decisions only.

Uncertainty. There is uncertainty over their next period wage draws ($\mathbf{q}_{t+1}^l \equiv \{\eta_{t+1}^m, \eta_{t+1}^f\}$) with distribution given by $F_{t+1}^l(\mathbf{q}_{t+1}^l)$ and there is now stochastic mortality (where assume that both members of the couple die in the same period).

State variables. The vector state variables ($\mathbf{X}^{1,l}$) during the late adult phase of life are the same as those for the early adult phase except that the (now-grown-up) child’s ability is no longer a state variable):

1. Age (t)
2. Assets (a)
3. Wage rates (w^m, w^f)
4. Education levels (ed^m, ed^f)
5. Own abilities (ab^m, ab^f)

⁷ Recall that the timing convention that we index all value functions in this exposition by the age of generation 1. That is, $V_{t+1}^{2,e}$ is the value function for generation 2 when generation 1 is aged $t + 1$.

Value function. The decision problem in the ‘late adult’ phase of life can be expressed as:

$$V_t^{1,l}(\mathbf{X}_t^{1,l}) = \max_{c_t^1, h^{m,1}, h^{f,1}} \left(u(c_t, l_t^m, l_t^f) + \beta s_{t+1} \int V_{t+1}^{1,l}(\mathbf{X}_{t+1}^{1,l}) dF_l(\mathbf{q}^l) \right) \quad (9)$$

s.t. the intertemporal budget constraint in equation (4)

and the time budget constraint in equation (5)

where s_{t+1} is the probability of surviving to period $t + 1$, conditional on having survived to period t .

4 Evidence on the returns to parental investments

In this section, we present regression analysis that relates parental time investments to children’s ability, and then in turn relates that ability to their adult earnings. This exploits a key advantage of our data — that we measure for the same individuals their parents’ investments, their ability, and the value of that ability in the labour market.

4.1 *The effect of time investments on ability*

In Section 2, we documented that the ability of children of higher-educated parents rises faster through childhood than the ability of other children, and that their parents spent more time reading to them and were more interested in their educational progress. We now look more formally at the relationship between those two facts using a simple regression framework.

To create a unidimensional measure of the time investments of parents in children (something that is required for the model outlined in Section 3 to be tractable) we extract a principal component factor from our proxies for the time investments of each parent, which is then normalised for ease of interpretation. At the age of 7, those proxies are frequency of reading with the child (for mother and father), teacher’s assessment of interest in education (for mother and father) and frequency of outings with the child (for mother and father). At the age of 11, the

proxies are library attendance, teacher’s assessment of interest in education and frequency of outings with the child (for mother and father).

4.1.1 Ability between 7 and 11

At the age of 11, we estimate the regression:

$$AB_{i,11} = \omega + \alpha AB_{i,7} + \beta^m TI_{i,7}^m + \beta^f TI_{i,7}^f + \delta ED_i + u_{i,11}^{AB} \quad (10)$$

where AB is normalised ability, TI^m is our (normalised) measure of mothers’ time investments, TI^f is our (normalised) measure of fathers’ time investments, and ED is parents’ education. This can be thought of as a simple approximation of equation 2, which governs the evolution of ability as the child ages (school type does not enter as it is constant across individuals until the age of 11).

The results from this regression equation are presented in Table 7. It shows that time investments have a significant effect on changes in ability over time, even after conditioning on background characteristics and initial ability. A one-standard deviation increase in maternal time investments at age 7 raises age-11 ability by 0.07 of a standard deviation, and a one-standard deviation increase in paternal time investments at age 7 raises age-11 ability by 0.04 of a standard deviation. These are relatively small effects, but are clearly statistically significant.

Table 7: Effect of time investments on the evolution of ability: 7 to 11

	(1)
	Normalised age-7 ability
Normalised age-11 ability	0.628 (0.00704)
Normalised age-7 maternal time investments	0.0739 (0.00845)
Normalised age-7 paternal time investments	0.0427 (0.00841)
N	10815

Standard errors in parentheses.

Regression includes controls for parental education.

Table 8 presents evidence of dynamic complementarity in the ability production function by adding to the estimating equation 10 interaction terms of child ability and time investments (see Agostinelli and Wiswall 2016). ‘Dynamic complementarity’ occurs when the return on present investments is higher if the initial level of the outcome variable (in this case ability) is higher, making past and present investments complements. We test for dynamic complementarity through a simple interaction of the return to parental time investments with the level of the child’s ability at the age of 7. The positive and statistically significant coefficients on these interaction terms indicates that the ability production function does, in fact, exhibit dynamic complementarity at this stage of childhood, with higher returns on time investments in children with higher initial ability (as found by Cunha et al. 2010).

Table 8: Dynamic complementarity in ability production function: 7 to 11

	(1)
	Normalised age-11 ability
Normalised age-7 maternal time investments	0.0769 (0.00848)
Normalised age-7 ability × Normalised age-7 maternal time investments	0.0176 (0.00825)
Normalised age-7 paternal time investments	0.0417 (0.00840)
Normalised age-7 ability × Normalised age-7 paternal time investments	0.0168 (0.00849)
N	10815

Standard errors in parentheses

Regression includes controls for parental education and age-7 ability

4.1.2 Ability between 11 and 16

At the age of 16, we estimate the regression:

$$AB_{i,16} = \omega + \alpha AB_{i,11} + \beta^m TI_{i,11}^m + \beta^f TI_{i,11}^f + \delta ED_i + \kappa st_i + u_{i,16}^{AB} \quad (11)$$

where the sole difference from the equation estimated at age 11 is the addition of the individual's school type st — comprehensive (the base category), grammar and private.

Table 9 shows the results from this regression. There are three key things to note. First, the coefficients on parental time investments are smaller than in the previous period, reflecting the fact that it is harder to affect ability between the ages of 11 and 16 than between 7 and 11 — as also shown by the larger coefficient on last period's ability. Second, the relative importance of mothers' and fathers' time investments are estimated to have reversed (though the difference between the two coefficients in the age-16 regression is not statistically significant). The point estimates suggest that an additional standard deviation of maternal time investments between 11 and 16 yield an increase of 0.03 of a standard deviation in age-16 ability, compared to an increase of 0.05 of a standard deviation for every additional standard deviation of paternal time investments. Third, school type is estimated to have a very significant impact on ability at 11

conditional on ability at 16. Relative to attending a comprehensive school, attending a private school is associated with an increase of 0.10 of a standard deviation in age-16 ability, and attending a grammar school is associated with an increase of 0.17 of a standard deviation.

Table 9: Effect of time investments on the evolution of ability: 7 to 11

	(1)
	Normalised age-16 ability
Normalised age-11 ability	0.760 (0.00740)
Normalised age-11 maternal time investments	0.0308 (0.0117)
Normalised age-11 paternal time investments	0.0452 (0.0117)
Grammar	0.177 (0.0197)
Private	0.101 (0.0289)
N	7302

Standard errors in parentheses.

Regression includes controls for parental education.

Table 10 shows the results when interactions of ability and time investments measures are included. These results suggest that, unlike investments between the ages of 7 and 11, there is not much evidence for dynamic complementarity between time investments between 11 and 16 and previous investments. In fact, the negative and statistically significant coefficient on the interaction between age-11 ability and subsequent paternal investments is evidence for the opposite of dynamic complementarity: During teenage years father's time investments have the biggest impact for children with less ability at the age of 11.

Table 10: Dynamic complementarity in ability production function: 11 to 16

	(1)
	Normalised age-16 ability
Normalised age-11 maternal time investments	0.0249 (0.0117)
Normalised age-11 ability × Normalised age-11 maternal time investments	0.00129 (0.0122)
Normalised age-11 paternal time investments	0.0530 (0.0118)
Normalised age-11 ability × Normalised age-11 paternal time investments	-0.0380 (0.0122)
N	7302

Standard errors in parentheses.

Regression includes controls for parental education, school type and age-11 ability.

4.2 *The effect of ability on earnings*

In the dynastic model with intergenerational altruism laid out in Section 3, parents do not receive any direct return from their children having higher ability at the age of 16. Instead, they include their children’s expected lifetime utility in their own value function, with a weight determined by the intergenerational altruism parameter λ . Hence parental investments in children’s ability (both through time and money investments in education) will be driven by the return to ability in the labour and marriage markets. Here we focus on the return to ability in the labour market, as measured by its impact on wages. We estimate the wage equation laid out in Section 3:

$$\ln w_t = \delta_0 + \delta_1 t + \delta_2 t^2 + \delta_3 t^3 + \delta_4 \ln ab_{16} + \delta_5 PT + v_t \quad (12)$$

for each gender and education group. Of course, ability has an important indirect impact on wages through its relationship with education, but it also has a direct impact on wages conditional on education. This is shown by Table 11, which plots the estimates of δ_4 for each gender and education group. The interpretation of these coefficients is that they are estimates of

the log-point increase in wages associated with a log-point increase in age-16 ability, conditional on education.

Table 11: Log-point change in earnings for a 1 log-point increase in ability

	Male	Female
High-school dropout	0.16	0.20
High-school graduate	0.31	0.29
Some college	0.55	0.38

The table shows that, as one would expect, age-16 ability has a significant positive impact on wages conditional on education for all groups. Perhaps more interesting, it finds evidence of complementarity between education and ability in the labour market, particularly for men. While male high-school dropouts see only a 0.15 log-point increase in hourly wages for every additional log-point of ability, men with some college education see an average increase of 0.55 log-points in hourly wages for every additional log-point of ability.

5 Results

In this section we present some findings on the quantitative importance of different investments and stages of childhood in explaining the intergenerational transmission of economic advantage. First, we conduct a simple 'back of the envelope' exercise in decomposing the difference in lifetime income between individuals from different parental backgrounds into the proportions explained by different channels of investment. This is limited in a number of ways discussed below, but provides powerful suggestive evidence about the sources of intergenerational transmission of inequality. Second, we use a simplified version of the model described in Section 3 to quantify the differences in expected lifetime utility by parental education, and to decompose those differences into the proportions explained by different

channels. This provides a more comprehensive measure of the relative importance of different channels, at the cost of relying on the structure of the model.

5.1 Decomposing the difference in lifetime income by parental education

In this analysis, we quantify the difference in expected lifetime income (as defined below) across children with fathers of the three different education levels defined and discussed in Section 2: compulsory, some post-compulsory and some college.

5.1.1 Methods

In this analysis, we focus on male members of the NCDS cohort, and define lifetime income as the sum of gross earnings during prime working age (between the ages of 25 and 55), plus any cash transfers and bequests from parents.

Differences in cash transfers and bequests from parents can be directly observed in the NCDS and ELSA data respectively, as reported in Table 6. To calculate differences in prime-age earnings we proceed in two steps. First, we estimate the earnings equation given in Sections 3 and 4. Second, we calculate in the NCDS data the distribution across education and ability levels of individuals with each level of father's education. By combining these two things, we can calculate expected lifetime earnings for each paternal education group.

Having calculated expected earnings for each paternal education group given the actual distributions of ability and education within each group, we then do the same calculation for three counterfactual distributions of ability and education across each paternal education group:

1. We predict the distribution of age-16 ability and education for each paternal education group conditional on age-7 ability. Differences in expected earnings across groups in this scenario reveal how much of observed differences in earnings

by paternal education can be explained by the differences in ability at age 7 shown in the first panel of Figure 1.

2. We predict the distribution of age-16 ability in the absence of differences in school type, and then predict education solely on the basis on that counterfactual age-16 ability distribution. The difference between expected earnings in this scenario and the previous one captures the effects of the faster growth in ability between 7 and 16 for children of higher-educated fathers, at least some of which is explained by the higher level of parental time investments in those children (as shown by the analysis in Section 4).
3. We use the actual distribution of age-16 ability, but predict the education distribution for each group on the basis of age-16 ability and school type, ignoring other factors. The difference between expected earnings in this scenario and the previous scenario captures the effects of schooling differences on lifetime earnings. The difference between expected earnings in this scenario and true expected earnings captures the effect on lifetime earnings of other drivers of educational outcomes besides ability and school type.

5.1.2 Results

Overall differences in expected lifetime income (as defined above) for men with different levels of paternal education are shown in the first row of Table 12. Those with mid-educated fathers have expected incomes more than £150,000 higher than those with low-educated fathers, and the gap between those with low- educated and high-educated fathers is almost exactly £300,000. For reference, the lifetime income of those with low-educated fathers is a little more than £850,000.

The rest of Table 12 decomposes these differences into distinct contributing factors.

- The first row of the decomposition shows differences in lifetime earnings in the first counterfactual scenario described above (age-16 ability and education predicted on the basis of age-7 ability). It shows that around 40% of the differences in lifetime income can be explained by differences in age-7 ability.
- The second row shows the difference between the first two counterfactual scenarios described above. It reveals faster growth in ability between 7 and 16 (not explained by different school types) explains around £23,000 of the gap between the children of low and mid-educated fathers, and around £39,000 of the gap between the children of low- and high-educated fathers (around 15% of the total gap in both cases).
- The third row shows the difference between the second and third counterfactual scenarios - schooling differences. The fact that those with higher-educated fathers are more likely to have attended private or grammar schools explains a little under 10% of the total differences in lifetime income
- The fourth row shows the difference between the final counterfactual scenario and actual expected earnings for each group. It suggests that differences in educational attainment conditional on ability and school type (explained by, for example, the role of financial support from parents) explains nearly 20% of the total gap in lifetime incomes across those from different parental backgrounds. It is perhaps surprising that differences in educational attainment conditional on school type and ability are twice as important in explaining differences in lifetime income as differences in school type.

- The final row of the Table simply documents differences in average inter vivos transfers and bequests across paternal education groups. It shows that around 20% of the differences in lifetime income across these groups are attributable to differences in transfers and bequests, rather than differences in earnings.

To summarise, the decomposition analysis suggests that around 40% of the difference in lifetime income across paternal education groups is attributable to differences in ability at age 7, around 40% by subsequent divergence in ability and different educational outcomes, and around 20% by inter vivos transfers and bequests received so far. Thus, while inter vivos transfers are important, most of the lifetime differences in lifetime income between children of low- versus high-education fathers are realized by the age of 16.

Table 12: Decomposition of differences in lifetime income by father's education

	Father's education	
	Some post-compulsory	Some college
Total difference	£156,000	£299,000
Accounted for by...		
Age-7 ability	£68,000	£115,000
Evolution of ability 7-16	£23,000	£39,000
School type differences	£11,000	£26,000
Attainment given ability and school type	£26,000	£58,000
Inter vivos transfers and bequests	£28,000	£61,000

Memo: Lifetime income for those with low-educated fathers: £854,000

Notes: Differences relative to those with low-educated fathers (compulsory education only).

Figures calculated for men.

5.2 *Decomposing the difference in expected lifetime utility by parental education*

There are a number of limitations with a comparison of expected lifetime incomes across individuals with different levels of parental education. Perhaps the most significant is that what individuals care about is the ex-ante difference in expected welfare, or expected utility. In this section we use a simplified version of the model laid out in Section 3 to estimate ex-ante

expected lifetime utility for children with each level of parental education, expressed using compensating variation.

5.2.1 *Methods*

The key simplification in the model used to estimate the results reported in this section is that we do not include intergenerational links - couples choose consumption and labour supply but not investments in their children. Hence the decision problem that households face corresponds to that described as the ‘late adult phase’ in Section 3. As a result, we do not explore how education and age-16 ability are determined within the model, but instead simply use the model to estimate expected lifetime utility given education and ability.⁸ We calibrate the preference parameters of this simplified model to match labour supply and wealth moments.

Our approach then roughly follows that described in the previous section. We first use the model to estimate expected lifetime utility for each level of education and ability, and then combine that with the distribution of individuals from each parental background across education and ability to estimate the actual expected lifetime utility for each level of father’s education. Then we can use the counterfactual distributions of education and ability discussed above to estimate expected lifetime utility for each parental education group in each of the counterfactual scenarios discussed. In order to provide a meaningful quantification of these differences in expected lifetime utility we calculate the consumption equivalent variation (CEV). This is the percentage increase in consumption in every state of the world required to make the children of

⁸ We also make a few further simplifications with respect to the model described in Section 3; namely, marital matching is on education only, we only allow individuals to choose whether to work full-time or not at all (no part-time choice), there are no earnings-related pensions (though each individual receives a flat rate pension in retirement) and preference parameters are not gender specific.

less-educated fathers indifferent between their ex-ante situation and that of those born to high-educated fathers.

5.2.2 Results

Table 13 shows the results from these CEV calculations. The first row of the table shows the total compensating variation required to make the children of low-educated fathers indifferent to being born to a mid-educated father (left-hand column) and a high-educated father (right-hand column). We estimate that the consumption of children of low-educated fathers in every state of the world would need to be increased by 6% for them to be indifferent with the children of mid-educated fathers, and by 12% for them to be indifferent with the children of high-educated fathers.

The rest of Table 13 decomposes these differences into distinct contributing factors.

- The first row of the decomposition shows differences in lifetime earnings in the first counterfactual scenario (age-16 ability and education predicted on the basis of age-7 ability). It shows that around 40% of the differences in expected lifetime utility (as measured by the CEV) can be explained by differences in age-7 ability. This is extremely similar to the proportion of the differences in lifetime income explained by age-7 ability in Table 12.
- The second row reveals faster growth in ability between 7 and 16 (not explained by different school types) explains between 15 and 20% of the differences in expected lifetime utility.
- The third row shows that schooling differences explain around 10% of the differences in expected lifetime utility - again a very similar result to that shown for expected lifetime income in Table 12.

- The fourth row shows that differences in educational attainment conditional on ability and school type (explained by, for example, the role of financial support from parents) explains around 30% of the total gap in expected lifetime utility across those from different parental backgrounds - a larger proportion than the 20% it explains of the difference in our measure of lifetime income. One potential reason for this is that the return to college attendance in the model is compounded by the existence of assortative matching in the marriage market - something not captured in our simple lifetime incomes analysis.
- The final row of the Table shows that inter vivos transfers explain only 2% of the difference in expected lifetime utility between those with low- and mid-educated fathers, and only 7% of the difference between low- and high-educated fathers. This is somewhat in contrast to the findings presented in Table 12, which show that inter vivos transfers and bequests explain around 20% of the differences across parental education in our measure of lifetime income. One reason for this difference is simply that bequests are not incorporated in the model, but perhaps more important is that these kinds of intergenerational transfers are extremely unequally distributed. Hence while they might have a meaningful impact on mean lifetime resources, from an ex-ante perspective they have very little effect on expected utility.

To summarise, this analysis of consumption equivalent variation (CEV) largely reinforces the conclusion of the previous analysis of expected lifetime income. Again, around 40% of the differences between those from different parental backgrounds can be explained by differences in ability by the age of 7, with the vast majority of the rest of the discrepancies being driven by

later differences in ability and educational attainment, rather than cash transfers from parents to children. In fact, this analysis reinforces the conclusion that most of the meaningful differences between children of low- versus high-education fathers are realized by the age of 16.

Table 13: Decomposition of differences in expected lifetime utility by father's education

	Father's education	
	Some post-compulsory	Some college
Total compensating variation	5.9%	12.3%
Accounted for by...		
Age-7 ability	2.6%	4.8%
Evolution of ability 7-16	1.2%	2.0%
School type differences	0.6%	1.3%
Attainment given ability and school type	1.6%	3.9%
Inter vivos transfers	0.1%	0.9%

6 Conclusions and policy implications

Understanding intergenerational links, and in particular the role of parental investments, is crucial for policymakers seeking to design redistributive tax and transfer policies that mitigate inequalities and improve social mobility, and wishing to understand the degree of intergenerational altruism (and hence the willingness of one generation to make sacrifices for another). In this paper we have documented substantial differences between children from different backgrounds in the evolution of cognitive ability through childhood, school quality and educational outcomes, and cash transfers received from their parents. A quantification of the implications of these differences for lifetime incomes suggests that around 40% of the gap between the sons of low- and high-educated fathers can be attributed to ability differences at the age of 7, a further 40% to subsequent differences in ability and educational attainment, and the final 20% to differences in the amount of inter vivos transfers and bequests received. The relative importance of these different stages of life and forms of investment is also found when using a simple calibrated model to estimate the welfare gains from better parents (as measured using

consumption equivalent variation). We provide evidence that at least some of the differences in ability and education are attributable to parental time investments in children and investments in school quality respectively.

All this has a number of implications for economic policy. At the most general level, the paper shows that policymakers interested in tackling the intergenerational transmission of inequalities need to consider policies designed to counter the inequality-increasing effects of each of the three forms of parental investment, since each proves to be quantitatively important in driving inequalities in income. Moreover, policymakers should bear in mind the substitutability of these different forms of investment - any attempt to shut down one channel of parental investments is likely to provoke a shift towards investment in other forms. In fact, the elasticity of substitution between these different forms of investment is a key determinant of the optimal policy response, and is something we will quantify in later versions of this paper.

A Appendix Table

Table 14: Proportion of children in each father's education group

	Father's education
Compulsory	75%
Post-compulsory	20%
Some college	5%

References

- Abbott, B., G. Gallipoli, C. Meghir, and G. Violante (2016). Education policy and intergenerational transfers in equilibrium. Technical report, Cowles Foundation Discussion paper.
- Agostinelli, F. and M. Wiswall (2016, July). Estimating the technology of children's skill formation. Working Paper 22442, National Bureau of Economic Research.
- Castaneda, A., J. Diaz-Gimenez, and J.-V. Rios-Rull (2003). Accounting for the U.S. earnings and wealth inequality. *Journal of Political Economy* 111 (4), 818–857.
- Caucutt, E. M. and L. Lochner (2012, October). Early and late human capital investments, borrowing constraints, and the family. Working Paper 18493, National Bureau of Economic Research.
- Charles, K. K. and E. Hurst (2003). The correlation of wealth across generations. *Journal of Political Economy* 111 (6), 1155–1182.
- Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity: The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics* 129 (4), 1553–1623.
- Cunha, F., J. J. Heckman, and S. M. Schennach (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica* 78 (3), 883–931.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies* 71 (3), pp. 743–768.

- Dearden, L., S. Machin, and H. Reed (1997). Intergenerational mobility in Britain. *The Economic Journal*, 47–66.
- Del Boca, D., C. Flinn, and M. Wiswall (2014). Household choices and child development. *The Review of Economic Studies* 81 (1), 137.
- Fuster, L., A. Imrohoroglu, and S. Imrohoroglu (2007). Elimination of Social Security in a Dynastic Framework. *Review of Economic Studies* 74 (1), 113–145.
- Gayle, G.-L., L. Golan, and M. Soytaş (2015). What is the source of the intergenerational correlation in earnings. Technical report, Mimeo.
- Lee, S. Y. and A. Seshadri (2016). On the intergenerational transmission of economic status. Unpublished manuscript, University of Wisconsin–Madison, Department of Economics .
- Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data. *The Review of Economics and Statistics* 87 (2), 235–255.