Explaining Income Inequality and Social Mobility: The Role of Fertility and Family Transfers.*

Diego Daruich† NYU

Julian Kozlowski‡ NYU

May 5, 2015

PRELIMINARY

Abstract

How much of social mobility and income inequality is due to initial opportunities relative to adult income risk? Previous studies have provided very wide estimates due to data limitations. To provide a more precise answer this article builds on a standard heterogeneous agent life cycle model with idiosyncratic income shocks. We propose that fertility differentials between rich and poor households can lead to substantial differences in the resources available for children, which can be important for their adult outcomes. Accounting for this is essential for the proper evaluation of initial opportunities so we extend the model to introduce the role of families through endogenous fertility, family transfers and education. We find that initial conditions as of age 13 account for more of adult income inequality than do labor income shocks. Moreover, fertility differentials and family transfers are found to account for over 50% of the social mobility in the data.

JEL Classifications: D91, J13, J24, J62.

Keywords: Fertility, Inequality, Social mobility, Quantitative model.

*We are grateful to our advisors Ricardo Lagos and Gianluca Violante. For helpful comments, we thank Raquel Fernandez, Violeta Gutkowski, Mi Luo, Virgiliu Midrigan, Simon Mongey, Debraj Ray and Edouard Schaal. We alone are responsible for all errors and interpretations.

†ddaruich@nyu.edu.

‡kozjuli@nyu.edu.

1 Introduction

Recent empirical work has increased the interest on social mobility and income inequality.[1] But is inequality mainly due to differences in opportunities determined early in life or to differences in effort and luck experienced over the working lifetime? What factors determine social mobility? And how are inequality and social mobility connected? Answering these questions is of utmost importance. First, it will help evaluate the different policies usually suggested to reduce inequality. Shall we provide insurance against shocks over the working lifetime (e.g. progressive taxation or unemployment insurance) or shall we focus on equalizing initial opportunities (e.g. early childhood investment or improving education access)? Similarly, an answer would help understand if those policies are effective in improving social mobility or if other measures may be necessary. Thirdly, theories of distributive justice typically distinguish ethically acceptable inequalities (e.g. due to differences in effort) from unfair inequalities (e.g. due to endowed characteristics) (Arneson, 1989; Cohen, 1989). Preferences for redistribution are systematically correlated with beliefs on the relative importance of these two in the determination of outcomes (Alesina and Giuliano, 2011). Those with beliefs that inequality is mainly due to differences in endowed characteristics tend to be more willing to accept redistributive policies. Hence any political discussion on redistributive policies is not likely to prosper without answer to these questions. However, providing one has been difficult in empirical applications due to lack of data so this paper uses an extended standard heterogeneous agent life cycle model with idiosyncratic risk to overcome this difficulty. Accounting for the role of families is essential for the proper evaluation of initial opportunities so we extend the model to allow for endogenous fertility, family transfers and education. We propose that fertility differentials between rich and poor households can lead to substantial differences in the resources available for children, which can be important for their adult outcomes. The model also allows for human capital transmission from parents to children, calibrated to match that in the data. We find that initial opportunities, as of age 13, account for more of the inequality observed in the data than do labor income shocks. Moreover, our results suggest that fertility differentials and family transfers generate over 50% of the (lack of) social mobility observed in the data.

Why is inequality a concern? Inequality can be thought of - and has been modeled - as a necessity to provide incentives for people to study or put effort (Krueger and Ludwig, 2015). However, Figure 1, based on Corak (2013) and known as the Great Gatsby Curve, which looks at income inequality and social immobility - the likelihood that someone will inherit their parents' relative position of income level - across (developed) countries has called public attention on one possible concern from income inequality. The Gatsby Curve shows that countries with high levels of inequality are associated with low levels of social mobility. One interpretation has been that (using a popular metaphor) as the rungs get further apart, it is harder to climb the lather. As inequality grows, more children may not be capable to improve their income relative to their parents. If children born from poor parents are not able to access education (e.g. due to borrowing limits or detrimental home environments) they may not be capable of taking advantage of the extra rewards or incentives from studying. This has been suggested to call for policy attention.

[1] See among others Atkinson et al. (2011), Aguiar and Bils (2011), Meyer and Sullivan (2013) or Chetty et al. (2014).

Figure 1: Gatsby Curve: Intergenerational Mobility and Income Inequality

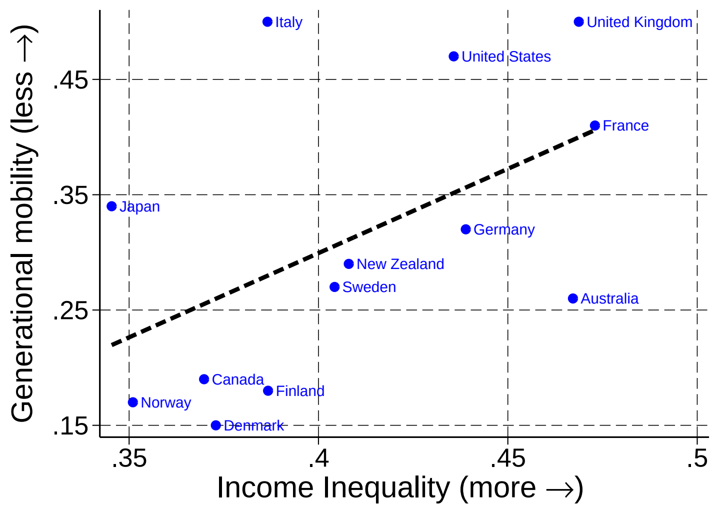

Note: Income inequality is measured using the Gini index on household income before taxes and transfers. Intergenerational economic mobility is measured as the elasticity between paternal earnings and a sons adult earnings. **Source**: Corak (2013) and OECD.

However, many interpretations are possible for Figure 1 so a proper framework to study income inequality and social mobility is necessary. Quadrini and Rios-Rull (2014) suggest that there is no conclusive theory to study inequality but that, among the economic literature, the classic Bewley economy is a typical one.[2] Probably the best example of a Bewley model used to study the sources of inequality is Huggett et al. (2011). However, they find that most of the income inequality is due to conditions given before entering the labor market which are exogenous in their analysis. More importantly, they admit that their results are silent about the forces prior to age 23, which can be considered very important when interested in inequalities of opportunities (or endowed characteristics). Previous empirical studies have tried to estimate the share of inequality that is explained by inequalities of opportunities but have been constrained by lack of data. The problem is that it is not possible to observe all initial conditions that determine opportunities. For example, direct data on the quality of the home environment or resources available for education are not usually available. This has forced these empirical studies to provide very wide estimates due to data limitations (Brunori et al., 2013; Niehues and Peichl, 2014). Their estimates (using the Theil-L index), suggest that between 16 and 75% of adult inequalities are explained by initial inequalities of opportunities. Though interesting, their wide range can turn them not useful for policy analysis. If we believe inequalities of opportunities are close to that lower bound, we might prefer redistributive policies that are focused on late redistribution. On the other hand, if we believe that inequalities of opportunities are closer to the upper bound we might find policies focused on improving the initial distribution of opportunities to be more appropriate. In order to overcome this data limitation, this paper introduces a model in the spirit of Huggett et al. (2011) but that endogenizes these earlier stages of life through choices of education, fertility and family transfers. As we push back the age at which initial conditions are determined, one's family becomes more important so including fertility and transfers decisions is essential. It would be interesting to go all the way back to age 0 so as to clear the mystery on how actual initial conditions are formed, but our model contributes by making one step in that direction and looking at initial conditions as of age 13. The calibrated model, which is consistent with the empirically observed relation between income inequality and social mobility from Figure 1, will provide a more precise answer to the impact of initial conditions.

Jones and Tertilt (2008) study fertility in the United States and provide substantial evidence that poor families

[2]Besides a typical Bewley economy, other theories that generate inequality include sorting (Ferndez and Rogerson, 2001; Ferndez et al., 2005) or models that can generate a Pareto distribution, like for example any with exponential growth occurring over an exponentially-distributed amount of time (Jones and Kim, 2014).

have more children than richer ones. In Section 2, we update their analysis for the US and extend it to 28 countries, showing the novel finding that this negative relation between fertility and income gets smaller as the development level of the countries increases, suggesting that fertility differentials narrow as countries experience economic growth. In their excellent literature review, Jones et al. (2010) mention that there is no full consensus on the motivations behind fertility choices but, to the best of our knowledge, the papers in the literature are not able to explain this finding. Most fertility models are very simple three-period representative agent ones which abstract from uncertainty issues.[3] This also makes them inappropriate for a quantitative analysis of the impact on income inequality or social mobility. We contribute to this literature by building a full life cycle heterogeneous agents model which allows for uninsurable shocks and is also capable of explaining the reduction in fertility differentials. Since we propose that fertility differentials and family transfers could be generating inequality and social immobility, capturing this finding is important if we want to understand which are the actual drivers of these family decisions.

Fertility differentials and family transfers can generate many effects on inequality and social mobility. Assume, as in most of the data, that poor families have more children than rich ones.[4] Having few resources, poor agents will be unable to transfer enough money for their children to afford higher education. While they have many children with few resources, rich households have few children with plenty of resources. These initial differences in resources can generate big differences in their adult outcomes through education. For the sake of clarity, abstract from family differences and take an economy where fertility differentials and family transfers do not exist. All children are born equal (with the same skills and resources) so they will all choose the same level of education. In such an economy the only source of inequality will be adult risk. Moreover, social mobility would be “perfect”: parents’ and children’s relative position of income would be independent. Now suppose you start from this economy and, all of a sudden, allow for endogenous fertility differentials and family transfers. Adults that were lucky and are currently rich will have few children and give them plenty of resources, while those on the opposite end will have many children with few resources. The first group of children will now be able to educate themselves even more than the previous generation, while the second group will not even be able to afford the previous generation’s level of education. As education increases their expected lifetime income, the children of rich households are more likely to remain rich. This will generate a reduction in social mobility. Moreover, income differences would also increase as the sources of inequality now include education access on top of adult risk. This is an extremely simple example but conveys the main mechanism through which fertility differentials and family transfers may affect inequality and social mobility.

The theoretical effect of fertility differentials and family transfers may be interesting in itself, but quantifying their impact is of utmost importance for policy questions regarding redistribution. Hence, we use our model to overcome data limitations and answer our initial questions. There are different ways to analyze the impact of initial conditions, but we broadly find that initial opportunities (as of age 13) account for more of the inequality observed in the data than do labor income shocks. More particularly, if we look at inequality of opportunities as computed by Brunori et al. (2013) or Niehues and Peichl (2014), our model suggests that 71% of observed inequality is due to inequality of opportunities (very close to their empirically estimated upper bound). On the other hand, if we focus on the variation of lifetime earnings, as does Huggett et al. (2011), our model suggests that 53% of this is due to differences in conditions determined as early as before high school. Similarly to them, these results should be understood as applying to the age where our model starts. Even though our model is silent about the forces determining the initial level of human capital (though calibrated such that the correlation between parents and children human capital holds as in the data), it can still shed light on its importance. More applied research (like Gertler et al. (2013)) may be needed to understand how this initial distribution may be improved. Finally, our model also suggests that fertility differentials and family transfers generate over 50% of

[3]See Morand (1999); de la Croix and Doepke (2003); and Mookherjee et al. (2012). Two exceptions might be Manuelli and Seshadri (2009) or **?** which move towards a full life cycle model to explain fertility differences across countries. Nevertheless, both of them abstain from uncertainty and, though heterogeneity is allowed in the second one, it is only in the form of constant skill differences across dynasties.

[4]For a clearer explanation of why this happens in our model see Section 3.

the social immobility observed in the data. Models interested in understanding social mobility should take these two forces into account.

The article is organized as follows. Section 2 shows the empirical work. Section 3 introduces the model, while Section 4 explains its estimation. Results are detailed in Section 5 and Section 6 concludes. The Appendix contains the list of countries used in the empirical work and some additional figures.

2 Empirical Findings

Since we are going to analyze social mobility and income inequality through the lens of a model with fertility decisions and family transfers, we first analyze the data available on these two. If children were considered a normal good, we should observe richer people having more children. However, this is not usually found in the data. On a country time series level, most countries have experienced a decrease in fertility over time (as they become richer). On a cross-country level, richer countries tend to have a smaller number of children. More importantly for our study, within a country-year it is also the case that richer people tend to have a lower number of children. Jones and Tertilt (2008) look at Census data for United States on women born between 1826 and 1960 and provide substantial evidence that the relationship between income and fertility was stably negative. Controlling for several factors (for instance, urban versus rural families, location or race), they suggest that economic factors play a big role in fertility decisions and that this relation with income is robust. We update their analysis for the US and extend it to 28 countries using micro data.[5] Note that around 40% of those countries are usually considered developed. We show that while the negative relation holds for most countries, it is also the case that it becomes less clear as the development level of the countries increases. In other words, it seems that as countries get richer, fertility differentials between the (relatively) poor and rich get smaller. Since our model will rely on old-age support motives to explain these observations, we also introduce evidence that old-age support is sizable in the data.

2.1 Fertility and Income

Economic models focus on the decisions made by individual households. Consequently, we would like a measure of fertility decisions at the household level. Probably the closest measure to this is available from the US Census: Children Ever Born (CEB). This variable asks each woman how many children they had had during their lives and allows researchers to compute fertility rates by cohorts. Unfortunately, this variable has some limitations. First, it requires women's fertility period to be over to be of use for our purposes. Even assuming that child bearing age extends only to forty years old, using the most current census possible only women born forty years ago could be used. Notice also that choosing the upper end of the age that determines the sample can bring issues. For example, if we used women up to any age we might get biased measures of fertility if this is correlated with mortality risk. Last but not least, this variable is only available for a very small set of countries and has even been dropped from the US Census after 1990. Hence, we use an alternative measure of fertility.

For the sake of clarity let us introduce the most basic measure of fertility, the Crude Birth Rate (CBR), which is defined as the ratio of births to women alive in one year. A typical issue with the CBR is that it can be too low because of a big share of women who have already completed their child bearing age, but are still pulling the ratio down. The Total Fertility Rate (TFR) attempts to correct some of these issues. It is defined as the

[5] For most of our empirical work we use census data from Minnesota Population Center (IPUMS) and household survey data from Luxembourg Income Study Database (LIS). Table A.1 in the Appendix details the countries and years used as well as the source for each case.

sum of the age-specific birth rates over all women alive in a given year. Hence, under the same example, if there is an unusually large number of women outside of the child bearing age, TFR is not affected. Formally, let $f_{a,c,t}$ be the number of children born to women of age a in country c and period t divided by the number of women of age a in country c and period t. Assume that the child bearing age extends between ages a_L and a_H.[6] Then the TFR in country c and period t, $\text{TFR}_{c,t}$, is defined as

$$\text{TFR}_{c,t} = \sum_{a=a_L}^{a=a_H} f_{a,c,t}.$$

Typically these age specific fertility rates are constructed for bands of ages of width 5 years and then summed, with the limits of the sum being $a_L = 15$ and $a_H = 49$.[7] Relative to CEB, the main benefit is that it does not require the surveys to report how many children has each woman had. Instead, it only needs for the children under the age of one to be associated to theirs mothers within the household - a much more standard requirement. Moreover, TFR does not require for the child bearing age to be complete as it focuses on fertility rates which are not associated with a particular cohort but with the women currently alive. Hence, information on the TFR is more up to date than that of the CEB. For this and other reasons, TFR has been widely used in the literature (Kremer and Chen, 2002; Manuelli and Seshadri, 2009).[8]

In order to connect the fertility rate with income, we define the TFR conditional on the income group. Suppose we divide the mothers according to their household income level in quantiles. Then, let $f_{a,q,c,t}$ be the number of children born to women of age a within quantile q in country c and period t divided by the number of women of age a and income quantile q in country c and in period t. Then, the TFR of income quantile q in country c and period t, $\text{TFR}_{q,c,t}$, is defined as

$$\text{TFR}_{q,c,t} = \sum_{a=15}^{a=49} f_{a,q,c,t}. \tag{1}$$

The appropriate measure of income is not obvious either. Assuming households have perfect foresight of their income, using their lifetime income would probably be the best measure. Jones and Tertilt (2008) use "Occupation Income" as their measure of choice. This is constructed for year 1950 by IPUMS and the authors extend it to their whole period of interest by assuming a constant 2% annual increase, equal across all occupations. This assumption does not seem harmless since occupations change their relative importance in the society over time (Autor et al., 2008). Moreover, there is a lot of variation in income across people within a same occupation.[9] Finally, this Occupation Income measure is not available for the panel study we are interested in. Hence, for this section we focus on annual total household income in the year of the sample. In order to get the appropriate quantile groups, we cannot compare the income level of young and old households since, following the typical life cycle of income, young households tend to have lower incomes. Hence, we define quantiles within the appropriate age group used for the TFR calculation.[10] This way the TFR for each quantile-country-year can be estimated.

[6]Notice that, assuming most women have children only in that period, extending this sample would most likely add only values of zeros to the formula of the TFR.

[7]Notice that when using age bands of width bigger than one year (but having only one year of data), $f_{a,t}$ is calculated as the number of children born to women within age band A in country c and in year t divided by the number of women within age band A in country c and in year t, *multiplied by the length of age band A*.

[8]The TFR measure of fertility also has its weaknesses. Since it is computed using data from a given year, it mixes fertility decisions of the different birth cohorts alive at the time. If all of these had the same fertility decisions, both CEB and TFR would be identical. However, if fertility rates are changing from cohort to cohort, then CEB gives the more accurate picture of fertility decision. Given the data limitations, we do our empirical work based on the TFR measure of fertility.

[9]For example, see the *National Compensation Survey: Occupational Wages in the United States, July 2004, Supplementary Tables* (Bureau of Labor Statistics, August 2005), p. 3; on the Internet at http://www.bls.gov/ncs/ocs/sp/ncbl0728.pdf (visited Jan. 21, 2015).

[10]For example, for households within the age group 15-19 years old, income quantiles are defined among other households in the same age group.

Let $\text{inc}_{q,c,t}$ be the median income of quantile q in country c and year t. Now we are ready combine the information on fertility and income.[11] We estimate

$$\ln\left(fert_{q,c,t}\right) = \alpha_{c,t} + \beta_{c,t}\ln\left(inc_{q,c,t}\right) + \epsilon_{q,c,t} \tag{2}$$

where q, c and t stand for the same elements as before. $\beta_{c,t}$ will be referred to as the *elasticity of fertility to income* for country c in year t. If this value is negative, richer households tend to have lower number of children. Values closer to zero imply that fertility rates are not related to income (at least according to the specification used). Finally, positive values of this elasticity imply that richer household tend to have more children than poorer ones. For the sake of clarity, Figure 2 shows the elasticity of fertility to income for the US (years 1960, 1970, 1980, 1990, 2000 and 2010) and Brazil (years 1970, 1980, 1990, 2000 and 2010). Each of this is computed using census micro-data available from IPUMS. Hence, millions of individuals are used to compute each of those points. On the vertical axis we have the values of the elasticity $\beta_{c,t}$ from (2). On the horizontal axis we place the logarithm of the real GDP per capita (PPP) of each country in each year, as a proxy for development. This is obtained from the Total Economy Database. Figure 2 shows that Brazil and US have different levels of fertility elasticities to income, as Brazil generally shows bigger fertility differentials than US. Moreover, it is also seen that as countries get richer this fertility differential seems to get smaller. This appears to be true both pooling countries and within countries over time, but we check this more formally below.

Figure 2: Income Elasticity of Fertility: Brazil and US

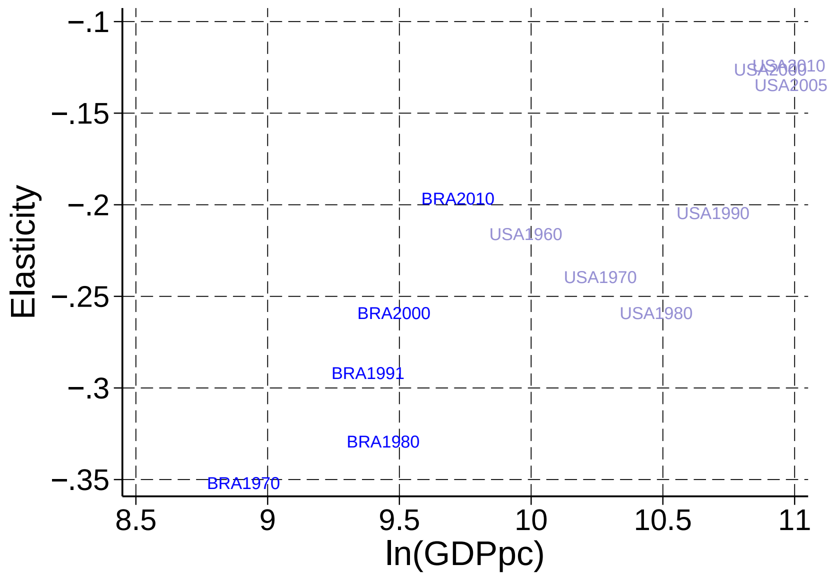

Source: IPUMS International. Methodology is explained in the main text. We remark that here income is computed at the household level since we do not have family identifiers in all the samples. It is possible, particularly among the poorer, that multiple families live within a household. We would like to separate their incomes between the families but are not able to for all the samples. However, when focusing on the US in the estimation section of the model, we can identify families and hence compute income at the family level (within the household).

We extend our analysis to include 28 countries for a total of 116 observations using micro data from both IPUMS and LIS. We combine all our observations and divide them into deciles according to their levels of real GDP per capita (PPP). Then, for each of these groups we calculate the median logarithm of the real GDP per Capita (PPP) and the median elasticity of fertility to income. Figure 3 shows that richer or more developed countries tend to have elasticities closer to zero or, in other words, smaller fertility differentials.

[11] Using mean income changes the results slightly, but they qualitatively remain the same.

Figure 3: Income Elasticity of Fertility by GDP Decile

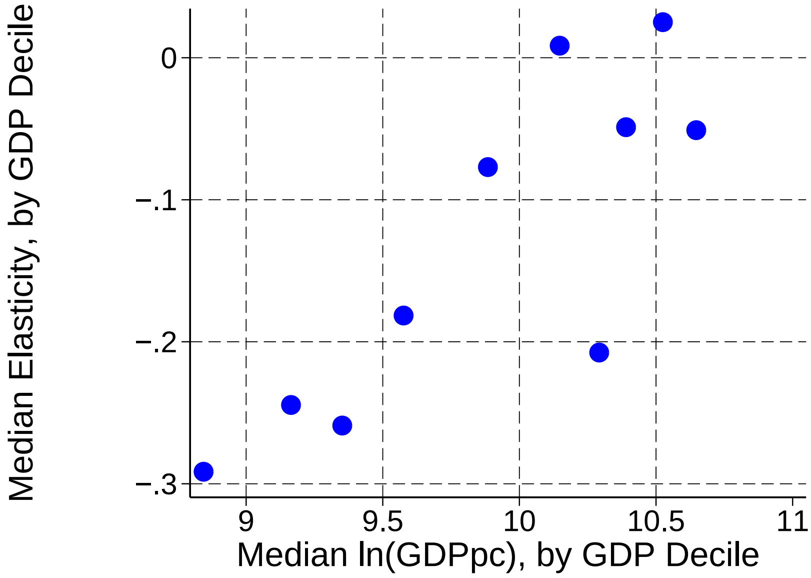

Source: IPUMS International and LIS. Methodology is explained in the main text. Figure A.1 in the Appendix includes all the data observations before grouping them into deciles.

Finally, various issues related to countries having different cultures or infrastructure, as well as recent years having better health resources or more female independence, might be hidden in those figures. In order to try to control for these possible issues we run a regression of fertility elasticities to income on the logarithm of the real GDP per Capita (PPP) controlling for either country or time fixed effects.[12] The regression specification is

$$\text{Fertility Elasticity}_{c,t} = \alpha + \beta \ln(GDP_{c,t}) + \eta_c + \mu_t + \epsilon_{c,t} \tag{3}$$

where Fertility Elasticity$_{c,t}$ is equal to $\beta_{c,t}$ from (2). Table 1 shows that the elasticity of fertility is increasing in the level of real GDP per Capita. Once again, this implies that more developed or richer countries are associated with smaller fertility differentials.[13] Moreover, this relationship seems stable and robust to controlling for country fixed effects as well as decade fixed effects. In Appendix Figure A.3 we also show that countries with more inequality are associated with higher fertility differentials than more equal ones. This evidence coincides with that shown by Kremer and Chen (2002) which looked at fertility differentials across education groups (while we focus on differentials across income groups).

[12]For time fixed effects we limit to decade fixed effects since very rarely do we have observations for different countries in the same year, as census or surveys take place in different times for each country. Moreover, we are not able to control for both country and time fixed effects since we do not have enough observations or variation in the data.

[13]It is important to remark that for some countries this relationship may actually be positive, meaning that richer households tend to have more children than poorer ones. This seems to be the case for countries were the population fertility rate is below the replacement ratio, like Finland or Denmark.

Table 1: Income Elasticity of Fertility and GDP per capita PPP

VARIABLES	(1)	(2)	(3)
ln(GDPpc)	0.207***	0.137***	0.203***
	(0.0423)	(0.0449)	(0.0424)
Observations	116	116	116
R-squared	0.160	0.871	0.167
# of countries	28	28	28
Country FE	NO	YES	NO
Decade FE	NO	NO	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Source: IPUMS International and LIS. Methodology is explained in the main text.

2.2 Family Transfers

Up to know we have presented evidence on two main facts: (i) Income inequality differs across countries and is associated with lower social mobility; and (ii) Higher levels of GDP per capita are associated with smaller fertility differentials. In Section 3 we introduce a life cycle model that is consistent with these patterns. Taking into account the importance in the development literature attached to old-age support when families choose the number of children as well as the magnitude of the family transfers documented, we include old-age support as a motive for fertility in the model.[14] Moreover, this will used in the model to capture the decrease in fertility differentials for richer countries. We now introduce evidence that old-age support - in the form of both money transfers and time - is sizable in the data.

Cox and Jimenez (1990) summarize the information on private transfers from 9 countries (including the US) and report that between 15 and 50% of people receive transfers annually. The higher end of that range is dominated by developing countries. Moreover, they also present evidence that, among developing countries, more than one-third of the elderly receive transfers from their children. But old-age support is relevant in terms other than money. For example, when health problems arise help from family can be essential. Lundberg and Pollak (2007) claim that in the US two-thirds of the 5.5 million elderly with disabilities rely on family for help.

We also use recent PSID information from the 2013 Rosters and Transfers survey to study transfers between parents and children. As suggested by Abbott et al. (2013), data on family transfers is scarce and problematic so more work needs to be done to extract more precise information from this type of sources.[15] Therefore, we take the evidence presented here as suggestive of mainly one fact: private transfers, with particular interest on those from children to parents, are substantial. For either direction of transfers, we calculate the average transfer over a lifetime.[16] The data is limited to children who are over eighteen years old and does not include "long-term"

[14] See Nugent (1985) or Banerjee et al. (2014).

[15] Estimates on the size of transfers depend substantially on whether observations with zero transfers are excluded or not.

[16] The procedure is similar to the one used to calculate the TFR in (1). We first calculate the average transfer given on the year before the survey (this is the question asked) by age groups. We then multiply this by the age width of the age group to obtain the average transfer given during that age window. Finally, we add up all age groups to obtain the average transfer over a lifetime. Different from Abbott et al. (2013) we also include observations with zero transfers. Note that we do not discount transfers to the present value.

transfers (for example, tuition or buying a house). In its first column Table 2 shows transfers going from parents to children, while the second column shows transfers going in the opposite direction. The first thing to notice is that these transfers are big, particularly once we include time provided to help either parents or children. Assigning them the mean wage value, we obtain that the average transfer from parents to children is almost 90% of the 2013 average annual household income. And transfers from children to parents are almost 60% bigger than those going in the opposite direction.

Table 2: Family transfers over a lifetime

	Parent → Children	**Children → Parent**
Money	\$38,589	\$37,536
Hours	1060	3166

Source: PSID Rosters and Transfers, 2013. Children are 18+ years old. It does not include transfers before that age neither those considered "long term" such as those for tuition or buying a house. Money and hours include cases of zero transfers.

3 Model

We specify a life cycle economy in a dynastic framework with three main stages. In the first stage, individuals make sequential education decisions: whether to acquire an extra year of education or start working. When education is complete, they enter the second stage, which represents their labor market experience. Idiosyncratic uninsurable income risk makes individual earnings stochastic. Throughout their life they choose saving and consumption expenditures. They can borrow only up to a limit, and save through a non state-contingent asset. During this stage, there is also a period when they choose how many children to have and how much resources to transfer to them. Finally, the last stage consists of retirement where they have three sources of income: savings, retirement benefits and old-age support from their children. We study the partial equilibrium version of this economy where prices and government policies are taken as given. We now describe the model and discuss the main mechanism.

3.1 The individual problem

Figure 4 shows the life cycle of an agent. Let j denote age. From $j = 1$ until $j = J_i$ the child lives with her parents who choose her consumption and at $j = J_i$ she becomes independent. The initial assets are money transfers from her parent and the initial human capital is stochastic but correlated with the human capital of her parent.

From $j = J_i$ until $j = J_e$ the agent has the option to study. The individual state variables are savings s_j and human capital h_j. While on education, she sequentially chooses whether to continue in school or to start working and this decision is irreversible. If she chooses education, her human capital increases deterministically by $f_j^e(h)$ and the cost of education is p_j^e.[17] While working, human capital evolves stochastically and is distributed by $f_j^w(h)$ where we allow for age-dependent idiosyncratic labor income shocks. In Section 4 we discuss the estimation of the returns of education and the income process.

[17] We allow education costs and returns to depend on the age of the agent in order to highlight the difference between high school and college in the estimation.

Figure 4: Life cycle

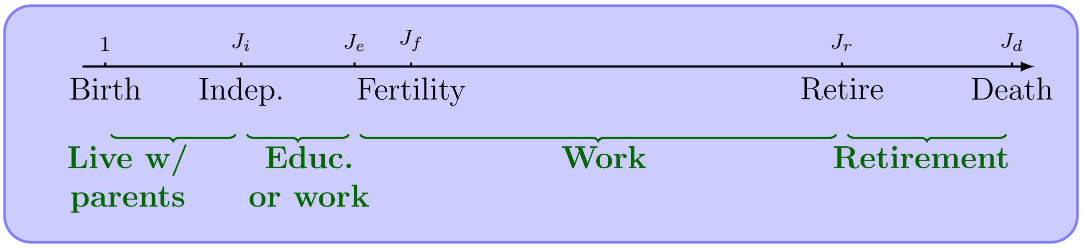

Formally, let V_j^e and V_j^w be the value of an agent of age j in school and working, respectively. Let V_j^{ew} be the value of an agent who can choose between the two,

$$V_j^{ew}(s_j, h_j) = \max\left\{V_j^e(s_j, h_j), V_j^w(s_j, h_j)\right\},$$

where V_j^e is defined by

$$\begin{aligned} V_j^e(s_j, h_j) = \max_{c_j, s_{j+1}} & \ u(c_j) + \beta V_{j+1}^{ew}(s_{j+1}, h_{j+1}) \\ & c_j + s_{j+1} + p_j^e = h w_j^e (1-\tau) + s_j (1+r) \\ & s_{j+1} \geq \underline{s}, \quad h_{j+1} = f_j^e(h_j). \end{aligned} \tag{4}$$

The agent is risk averse and her preferences are represented by the increasing, concave and positive utility function u.[18] She can borrow up to the limit $\underline{s}$ and the return on savings is $1+r$. Future is discounted by β and note that, in this particular problem, there is no uncertainty. We denote as w_j^e the wage for an agent of age j that is in school. In particular, we assume that the agent does not work during high school ($w_j^e = 0$) and we allow for (part-time or internship) work while in college.

The value of work V_j^w is defined by

$$\begin{aligned} V_j^w(s_j, h_j) = \max_{c_j, s_{j+1}} & \ u(c_j) + \beta \mathbb{E}\left[V_{j+1}^w(s_{j+1}, h_{j+1})\right], \\ & c_j + s_{j+1} = h_j w (1-\tau) + s_j (1+r), \\ & s_{j+1} \geq \underline{s}, \quad h_{j+1} \sim f_j^w(h_j). \end{aligned} \tag{5}$$

The return from working is the wage w net of taxes τ. Note that there is no disutility from working and so the labor supply is inelastic. Also note that the choice of leaving education is irreversible. Once the agent enters the labor force, she cannot return to school.

From $j = J_e$ until $j = J_r$ the agent works and her individual problem is equivalent to (5). There are two special periods where the agent problem will be different and the number of state variables will change from then on. First, in the exogenously given fertility period $j = J_f$, the agent chooses the number of children and the money transfered to them. Second, when her parents retire at $j = J_t$, the agent provides old-age support, transferring a fraction of her current labor income to her parent.

We model altruism à la Barro and Becker (1989) where parents care about the utility of their children. Then,

[18]That u is positive is important to model altruism. As shown by Jones and Schoonbroodt (2009) the implicit assumption that parents like having children requires that the utility function must be always positive or always negative. If we choose the negative case we need an extra assumption for the value of having zero children. Therefore we follow the classic approach of u being always positive and assume that having zero children generates zero utility.

the problem at the age of fertility is

$$V_j(s_j, h_j) = \max_{c_j, c_j^c, s_{j+1}, n, \varphi, T} u(c_j) + \beta \mathbb{E}\left[V_{j+1}(s_{j+1}, h_{j+1}; n, \varphi, h_j)\right] \tag{6}$$

$$+ b(n)\left\{u(c_j^c) + \beta^{J_I}\mathbb{E}\left[V_{J_i}(\varphi, h^c)\right]\right\},$$

$$c_j + nc_j^c + s_{j+1} + \frac{n\varphi}{(1+r)^{J_i}} + C(h, n, T) = h_j w(1-\tau) + s_j(1+r),$$

$$s_{j+1} \geq \underline{s}, \quad h_{j+1} \sim f_j^w(h_j), \quad h^c \sim f^c(h_j).$$

In this period the agent chooses her consumption c_j, her children consumption c_j^c, savings s_{j+1}, the number of children n and the transfer to each children φ. This transfers are assumed to be equal for all children.[19] As usual, the agent derives utility from her own consumption and her continuation utility, whose states are explained below. Furthermore, similar to Roys and Seshadri (2014), the agent is altruistic and derives utility from her children consumption and the continuation utility of her children when they become independent. The altruistic discount factor $b(n)$ is increasing and concave.

Raising children is costly. The parent pays the cost $C(h_j, n, T)$ on top of the money spent on children's consumption and transfers. This cost is assumed to be increasing in the number of children n and in the level of human capital of the parent h_j. The variable T is an indicator variable, $T \in \{\text{hire childcare}, \text{stay at home}\}$, to capture the idea that childcare is a good that can be either produced at home or hired in the market. Parents with higher human capital will optimally choose to hire childcare in the market whereas parents with lower human capital will choose to stay at home and pay this costs with her own human capital.[20]

After the fertility decision the individual problem is similar to (5) with the following caveats. First, for J_i periods the parent chooses the children's consumption and pays cost C. Later, at $j = J_t$ a fraction of the current labor income goes to the agent's parent as old-age support. There are different ways to introduce old-age support. Altig and Davis (1993) allow for double-sided altruism (children taking into account the utility of their parents as well as parents taking into account that of their children) in a 3 periods model without heterogeneity. Even in this much simpler model, double-sided altruism brings many difficulties which lead them to eliminate linkages and strategic behavior from the agents. More recent attempts to include only the altruism on the children side have also been limited to representative agents economies without uncertainty (Boldrin et al., 2005). In our heterogeneous agent model, which includes multiple stages of discrete choices (education and fertility) as well as high dimensionality of the state space, the altruism approach to old-age support is computationally infeasible. Hence, we adopt the rule that children are constrained to transfer an exogenous share ξ of their income to their parents à la Morand (1999). Anyway, we believe that endogenous old-age support should actually strengthen our proposed channel, since - abstaining from strategic behavior issues - poor parents would expect to receive relatively higher transfers from their children as their needs are much more pressing than those from richer parents.[21] This would make the old-age support motive even more important for poor families.

At $j = J_r$ the agent retires with three sources of income. He has savings and retirement benefits that depend on the human capital and are progressive as in the US Social Security Administration. Furthermore, at the first period of retirement the agent receives transfers from her children as old-age support. Parents need to predict

[19] The altruism value derived from children depends on their initial assets. Therefore, we assume that parents set a fund at the age of fertility such that their children receive φ when they become independent. Moreover, we note that the utility and value functions are specified at the household level (with two individuals). Hence, for example, when parents choose to have two children, they create one household so $n = 1$.

[20] We allow for childcare for quantitative reasons. We explored different costs functions and we found that a model that allows for childcare fits the data on the elasticity of fertility the best. Without childcare, fertility differentials are too big relative to the data and do not get smaller with higher levels of income as reported in Figure 3. Finally, note that child care has no effect on the children's human capital.

[21] The marginal utility of consumption for poor parents would be higher than for richer families, increasing the incentive of children of poor parents to transfer.

how much money they will receive from their children when they get old. An extreme view is that parents know their children income perfectly, updating it year by year. Independently of the plausibility of this view, in our model this would require to extend the state space to include that of each children.[22] Added to the current dimensionality of the model such a procedure would become computationally infeasible. Hence, we assume that parents have limited information to predict the transfers they will receive. The only information that parents have about their children are the number of children n, the initial assets of their children φ and their own human capital at the age of fertility h_{J_f}.[23] These state variables remain constant until retirement and are used to predict the old-age support that the agent will receive. Note that the old-age support is an endogenous random variable whose distribution depends on the education choice of the children.[24] Formally, the problem at the age of retirement is

$$
\begin{aligned}
V_j\left(s_j, h, \theta\right) = \max_{c_j, s_{j+1}} & \; u\left(c_j\right) + \beta V^w_{j+1}\left(s_{j+1}, h\right), \\
& c_j + s_{j+1} = \theta + \pi(h) + s_j\left(1+r\right), \\
& s_{j+1} \geq \underline{s},
\end{aligned}
\tag{7}
$$

where θ are the old-age support transfers and π are the retirement benefits, which depend on the human capital at the age of retirement.

3.2 Fertility choices

Two incentives for having children are at play: (i) Altruism, from which parents care about their children's well-being; and (ii) Old-age support, from which parents care about the help their children provide them once they grow old. Altruism implies that parents want to have educated children (as this brings their children more income). In the calibrated version of the model, altruism will be such that it provides incentives to have a rather constant number of children across income groups. Old-age support causes children to be a form of private investment, with different rates of returns across families, which can be important as an investment for retirement. However, in general they are a relatively bad one since they are costly in a stage of life where they would actually like to borrow.[25] More importantly, for high income parents children are particularly costly due to the time they require. And their return in future transfers are also not that relevant since they will have plenty other sources of savings. The opposite is true for low income households, who have low time costs and obtain high marginal returns from them, making the old-age support channel much more valuable for poorer households. This way, the altruism channel dominates for richer households while the investment one gives an extra incentive for poorer agents to have children.

To understand these effects consider the first order condition with respect to the quantity of children n[26]

$$
u'_{J_c}\left[\varphi + C_n\right] = b_n\left(n\right)\beta^{J_I}\mathbb{E}_{J_c}\left[V_0\left(h^c_0, \varphi\right)|h_{J_c}\right] + \beta^{J_R - J_c}\frac{\partial \mathbb{E}_{J_c}\left[V_{J_R}\left(h, s, \theta\right)|n, \varphi, h_{J_c}\right]}{\partial n}. \tag{8}
$$

[22]Given that the data shows many families having over four children, this number of states is not small.

[23]The human capital at the age of fertility h_{J_f} is essential for the initial draw of human capital of the children. The initial assets are very important to determine how long the children will remain in school. Finally, the number of children n is important to determine, among others, the size and uncertainty of the transfer to be received.

[24]We remark that the difference with respect to the full information case is such that uncertainty on children's transfers last longer, but at the time of fertility the information available is equal. Hence, in a model with exogenous labor supply like ours, the assumption should only affect savings. As our focus is on labor income inequality, our assumption does not seem very harmful.

[25]For example, in the US the current average age of first birth is 27 while the income peak is closer to 50 years old.

[26]To simplify the exposition we do not take into account the utility derived by the consumption of the children. This argument is standard but complicates the explanation. A similar argument can be derived from the first order condition with respect to the transfers to children φ.

The left hand side is the marginal cost of an extra child, that is the transfer φ and the marginal childcare cost C_n, scaled by the parent's marginal utility at age of fertility u'_{J_c}. The right hand side is composed of two terms. The first is the benefit from altruism, while the second is from old-age support. The altruism channel is standard, with the payoff being the expected value function of the child and the discount factor increasing by $b_n(n)$. The benefit of old-age support is generated by the change in the distribution of the transfers θ. For the sake of clarity, we can decompose the random variable θ into its conditional mean and a martingale shock. This implies that

$$\begin{aligned}\frac{\partial \mathbb{E}_{J_c}\left[V_{J_R}|n,\varphi,h_{J_c}\right]}{\partial n} &= w\left(1-\tau\right)\xi\mathbb{E}_{J_c}\left[h^c_{J_T}|\varphi,h_{J_C}\right]\mathbb{E}_{J_c}\left[u'_{J_R}\right] \\ &+ \int\left(\int u'_{J_R}\frac{\partial g_\varepsilon\left(\varepsilon|n,\varphi,h_{J_C}\right)}{\partial n}d\varepsilon\right)df\left(h_{J_R}|h_{J_C}\right)\end{aligned}$$

where the expected transfer of each child is $w\left(1-\tau\right)\xi\mathbb{E}\left[h^c_{J_T}|\varphi,h_{J_C}\right]$. The first term is the effect on the conditional mean and shows that as n increases, so does the expected transfer. Once again, note that this is scaled by the expected marginal utility at the age of retirement. The second term reflects the higher order moments effects.

With this first order condition we can learn why for richer agents the old-age support is less important and fertility choices are dominated by the altruism channel. Note that both the marginal cost and the old-age support benefits are scaled by the marginal utility of consumption. Moving towards richer agents, the marginal utility diminishes and both the cost and the old-age support are less important. On the other hand, the altruism benefits do not decrease and therefore dominate the fertility trade-off. Consequently, this implies that as the economy grows - for instance as w increases - fertility choices will be controlled by altruism instead of old-age support. This provides a framework that generates the non-homothetic relationship between development and fertility choices that we documented in Section 2. In the calibrated model, altruism motives will always be present and lead families to have few children. This will be the main fertility driver for developed economies in general as well as for rich agents in poorer ones. On top of the altruism motive, poorer economies or low income agents will take into account the old-age support channel, leading them to have more children and generating the negative fertility elasticity. Moreover, the return of the investment on old-age support differs among families due to mean reversion of income across generations. On the one hand, poor families expect to have relatively richer children and therefore transfers will be larger. On the other hand, richer families expect to have relatively poorer children with low transfers. Hence the expected return of old-age support is larger for poor families which reinforces the negative income-fertility relationship.[27]

We can evaluate the effects of old-age support under the benchmark calibration described in Section 4. At the age of retirement there are three sources of income: social security, savings and old-age support transfers. Figure 5 shows the average contribution of each source across income quintiles.[28] For the poorest quintile old-age support represents over 40% of resources but for the richest quintile it represents less than 10%. There are two forces at play. First, due to the negative income-fertility relationship poor households choose to have more children and therefore receive larger transfers. Secondly, richer households accumulate more savings so the transfers from the children are less important. This coincides with the intuition obtained from equation (8): old-age support is more relevant for poor agents in our model.

[27]Once an economy is sufficiently rich, it is possible for children to behave as a normal good. Every household prefers educated children and rich families can afford more of them.

[28]Recall that old-age support transfers occurs only at age J_r but social security benefits are received throughout retirement. In Figure 5 we compare the net present value of social security with the stock of savings and the old-age support transfers.

Figure 5: Budget's Assets Distribution by Income Group

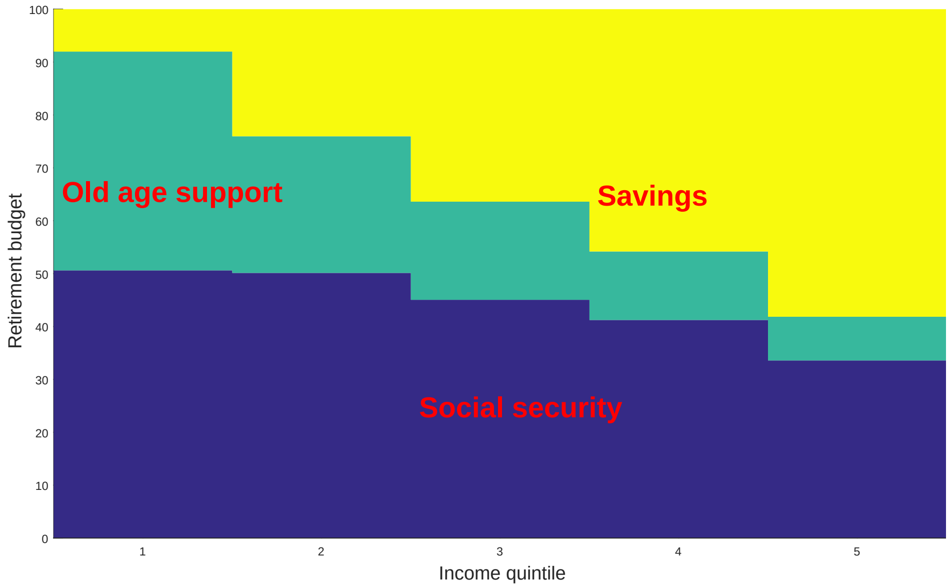

Note: Average income contribution of social security, savings and old-age support for each income quintile. Results depend on the estimation of the model, shown in Section 4.

4 Estimation

We numerically solve the steady state of this economy. Due to the presence of nonlinearities and discrete choices we implement a global solution. Some of the computational challenges are that we have up to five state variables and several non convexities due to the discrete choices in education and fertility. Therefore we apply a generalized endogenous grid method from Fella (2014). Then we simulate the economy to match moments from United States in 1960. Some of the parameters can be estimated "externally", while others need to be estimated "internally" from the simulation of the model. We first describe those in the first group.

Demographics: A period in the model is one year. Individuals become independent at the age of $J_i = 12$ and they start with the equivalent of 7 years of education. They can go to high school (five years) and then to college (four years) and so the maximum age for education is $J_e = 22$. Fertility decisions are made at the average age of first birth, $J_f = 27$. Retirement occurs at $J_r = 65$ and therefore parents retire at the child's age of $J_t = 38$, which is when children transfer money to their parents. Death is assumed to occur for all agents at age $J_d = 80$.

Prices: We normalize the wage to $w = 1$ and we estimate the wage while in college from IPUMS Census data. We focus on individuals between the ages of 18 and 22 years old and match the relative earnings of those currently in college relative to those who are not, leading to $w^e = 0.56$. Following Roys and Seshadri (2014) we set the interest rate to $r = 5.5\%$. We assume borrowing is not possible, $\underline{s} = 0$. The labor tax is from Mendoza et al. (1994), $\tau = 0.28$. The price of college is from Delta Cost Project, where we get $p_j^e = \$799$ for $j > J_i + 5$.[29] In this class of models it is difficult to match the high school dropouts rate. Previous studies such as Abbott et al. (2013) introduced nonpecuniary (psychic) costs of education. Instead, in this paper we estimate the price of high school to match the dropouts rate. Our estimates are reported in Table 5 with a price of high school of \$87. Note that our estimate of high school cost is about 10% of college cost which is consistent with the US

[29]We take into account grants and scholarships, such that only private tuition costs are considered. Prices are in 1960 dollars.

education system (relatively low cost of high school when compared to college).

Education returns: We base our estimates of the education returns on Table 3.a from Heckman et al. (2006). The yearly return for high school and college are 17% and 12% respectively. The degree completion premium for high school and college are 16% and 10% respectively.

Labor income risk: We assume that $f_j^w(h) = h(1+\delta)$ with $\delta \in \{\delta_u, \delta_d\}$ independent of age. However, we allow for the probabilities of the shocks to be age-dependent. We assume that Prob($\delta = \delta_u$ at age j) $= p(j) = p_0 e^{\gamma_h j}$ for $j \geq J_e$ and $p(j) = p(J_e + 1)$ is constant before the education stage (for $j < J_e$) as data is limited for those ages. We estimate δ_u, δ_d, p_0 and γ_h to match the first difference of mean and variance of log earnings at age 25-60. Two comments on this are appropriate. First, this income risk is calibrated to include total earnings variation, encompassing what may be considered both wage shocks as well as hours worked (or effort) differences. Second, even though we propose a simple model of income, we are able to match standard statistics of labor earnings. This is very important in order to properly evaluate the impact of initial opportunities on income inequality. Otherwise the comparison could be favorable for initial opportunities. Figure 6a shows the average income age profile, which is seen to be very well estimated. Figure 6b shows the variance of log income for those between 27 and 60 years old, always in differences relative to those aged 26. This shows the if anything we overestimate the growth in variance of the process. Together with our target below on the average income and variance level at age 25, this suggests that our estimates lay on the conservative side about the importance of inequality of opportunity relative to labor income risk. Finally, Figure 6c and 6d show two measures of income inequality over the life cycle, the ratio of top 20% to bottom 20% earners and the Gini coefficient. It is seen that both of the them display a similar trend both in the model and in the data.

Figure 6: Adult income risk

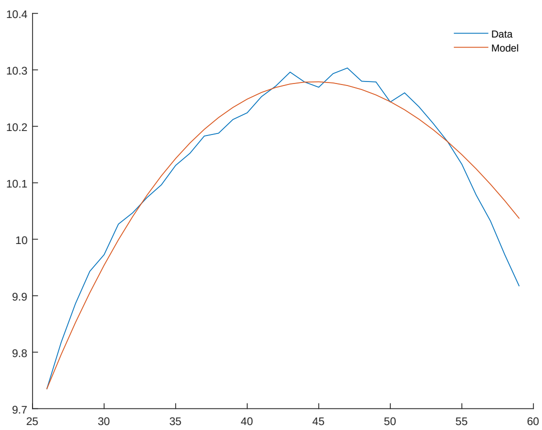

(a) Life cycle of log(income)

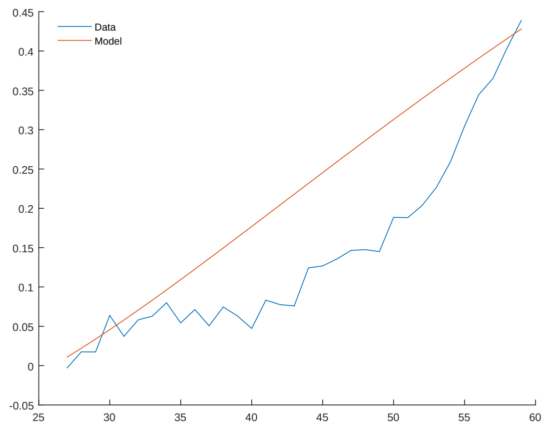

(b) Variance difference wrt age 26

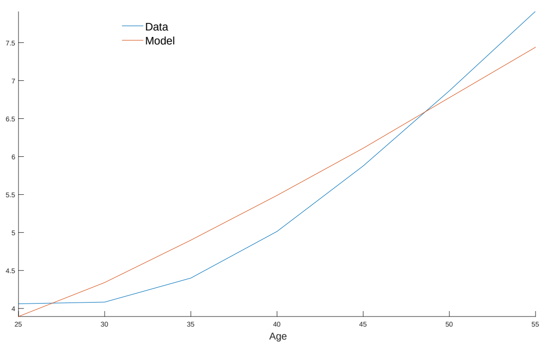

(c) Top-bottom

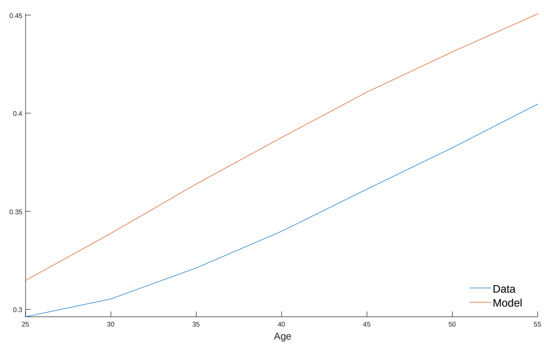

(d) Gini

Childcare: We propose the following functional form for childcare

$$C(h,n,T)=\begin{cases} n^{\alpha_n}\left(wh(1-\tau)\beta_c+(\beta_n-\beta_c)\,p_c\right) & \text{hire childcare} \\ n^{\alpha_n}\left(wh(1-\tau)\beta_n\right) & \text{stay at home} \end{cases}.$$

This function allows for non-constant returns to scale in the number of children. We particularly assume there increasing returns to scale with $\alpha_n = 0.7$.[30] Moreover, this functional form also allows for a reduction in the time needed to take care of children if childcare is hired in the market, which is done by around 35% of the population.[31] We estimate that if the agent stays at home, her labor income is reduced by $\beta_n = 1/6$ but if she hires childcare her labor income is reduced only by $\beta_c = 1/16$.[32] The cost of childcare is $(\beta_n - \beta_c)\,p_c$ with $p_c = \$2732$.[33]

Old-age support: With data from PSID Rosters and Transfers 2013 we back out the average net transfers to parents over the lifetime (after the age of 18 and not including schooling costs). In the model old-age support occurs only in one year, so we input all these transfers as if they where given only at age J_t. Based on Table 2, we estimate that the fraction of income that goes to parents as old-age support is equivalent to 59%.

Replacement benefits: The pension replacement rate is based on the Old Age Insurance of the US Social Security System. With the last level of human capital before retirement we estimate the average life time income to be $\widehat{y}\,(h) = 1.44 \times h$. Then, the pension formula is given by

$$\pi(h)=\begin{cases} 0.9\widehat{y}\,(h) & \text{if } \widehat{y}\,(h)\le 0.3\bar{y} \\ 0.9\,(0.3\bar{y})+0.32\,(\widehat{y}\,(h)-0.3\bar{y}) & \text{if } 0.3\bar{y}\le\widehat{y}\,(h)\le 2\bar{y} \\ 0.9\,(0.3\bar{y})+0.32\,(2-0.3)\,\bar{y}+0.15\,(\widehat{y}\,(h)-2\bar{y}) & \text{if } 2\bar{y}\le\widehat{y}\,(h)\le 4.1\bar{y} \\ 0.9\,(0.3\bar{y})+0.32\,(2-0.3)\,\bar{y}+0.15\,(4.1-2)\,\bar{y} & \text{if } 4.1\bar{y}\le\widehat{y}\,(h) \end{cases}$$

where $\bar{y} = \$5500$.

Intergenerational transmission of ability: We assume that the initial level of human capital is stochastic but correlated with the human capital of the parent. From NLSY79 we back out the transition matrix between parents' income and children's ability (measured according to the reading test) by quintile, shown in Table 3. First, conditional on their parents' income quintile, each child draws her ability quintile from Table 3. Then, she draws h_{J_i} uniformly within the given quintile of a log normal distribution with parameters (μ, σ) to be estimated.

Table 3: Correlation of parent's income and children ability, by quintile

	Children					
Mothers	1	2	3	4	5	Total
1	41.31	26.58	13.91	10.53	7.67	100
2	26.36	26.86	15.96	17.54	13.28	100
3	17.78	24.2	18.95	21.67	17.4	100
4	11.01	21.93	19.86	24.98	22.22	100
5	5.95	15.02	18.15	28.33	32.56	100

Source: NLSY79. Each cell reports the conditional probability in %.

[30]The returns to scale are based on the NAS equivalent from Michael and Citro (1995).

[31]Based on NLSY79

[32]We adapt the estimation from Angrist and Evans (1998) to be consistent with the equilibrium of our model. In particular we assume that in the estimation of Angrist and Evans (1998) high income households hire childcare in the market whereas low income households stay at home.

[33]The cost of childcare is based on the average wage of nannies from IPUMS Census data.

Preferences: We specify the period utility over consumption as a CRRA function

$$u(c) = \frac{c^{1-\gamma_c}}{1-\gamma_c}.$$

As discussed in Section 3 the utility function has to be positive and therefore $\gamma_c \in [0,1)$. We follow the literature and assume $\gamma_c = 0.5$ (for example, see **?**).

We are left with five parameters. Two parameters, λ_n and γ_n, are related to altruism. Other two, μ and σ, relate to the initial distribution of human capital while the last one is related to the high school cost. In order to pin down the value of these parameters we aim to match nine moments from the data, shown in Table 4. First, to have consistent prices as well as a reasonable income process, the mean and coefficient of variation of income at age 25 is targeted. Recall that the income age profile relative to this age is matched externally. Next, to identify the altruism we target the mean fertility and the fertility elasticity. Given our focus on inequality, we target the distribution of education attendance (dropouts, high school graduates and college graduates) and the top-bottom income ratio. Finally, as we are also interested in social mobility, we target the intergenerational mobility rank-rank coefficient.[34] Table 4 shows the results of the estimation, while table 5 shows the estimated parameters.

Table 4: Targeted moments

Moment	Data	Model
Mean income age 25	\$4337	\$4557
CV income age 25	0.086	0.069
Dropouts	34.1%	19.7%
High school graduates	52.8%	67.8%
College graduates	13.1%	12.5%
Mean fertility	3.400	3.666
Fertility elasticity	-0.150	-0.151
Intergenerational Mobility: Rank-Rank	0.341	0.440
Income top-bottom 25-59	5.286	5.337

Source: Mean and CV income at age 25, fertility elasticity and income top-bottom 25-29 are constructed from IPUMS Census data. Mean fertility is obtained from United Nations Statistics. Education attainment are reported in the Current Population Survey historical time series tables. The intergenerational mobility is from Chetty et al. (2014).

Table 5: Estimated parameters

μ	σ	γ_n	λ_n	$HScost$
-2.183	0.310	0.448	0.109	\$87

The model performs generally well. As explained above, given that we are working in a partial equilibrium framework matching the mean income is essential to have sensible prices. The estimation succeeds in this dimension. Secondly, fertility levels and elasticity are also successfully matched, which is very important given

[34] Given the very different values of the moments, we adjust the weighting matrix such that the difference used in the objective function of the minimization is closer to the percentage deviation. More precisely, moment j is weighted by $\frac{1}{\max\{\underline{m}, |m_j|\}}$ where m_j is the value of moment j in the data. Since some moments could potentially have the value of 0, we limit the denominator to be above $\underline{m} = 0.1$. We have tried other alternatives to solve for this issue, but we have not found significant differences.

the key role of fertility in our model. Regarding education, college graduates are well matched but matching high school dropouts is harder. This problem has been encountered by other researchers, many of whom have opted for adding exogenous psychic costs. However, given the lack of substantial evidence on what this cost actually represents, we avoid adding them to our model. With regards to inequality, the model displays similar levels of lifetime inequality as the data. We obtain relatively low levels of social mobility, but we can partially open the transition matrix on which this estimate is based on to get a better sense of why this is the case. Figure 7 shows that the model does a good job in replicating the transition probabilities for all entries except for the probability of remaining in the bottom quintile, where the model overestimates this probability. This is also why when we estimate the log-log coefficient as an alternative measure of social mobility the model is very close to the data (0.334 in the model and 0.344 in the data).

Figure 7: Intergenerational mobility

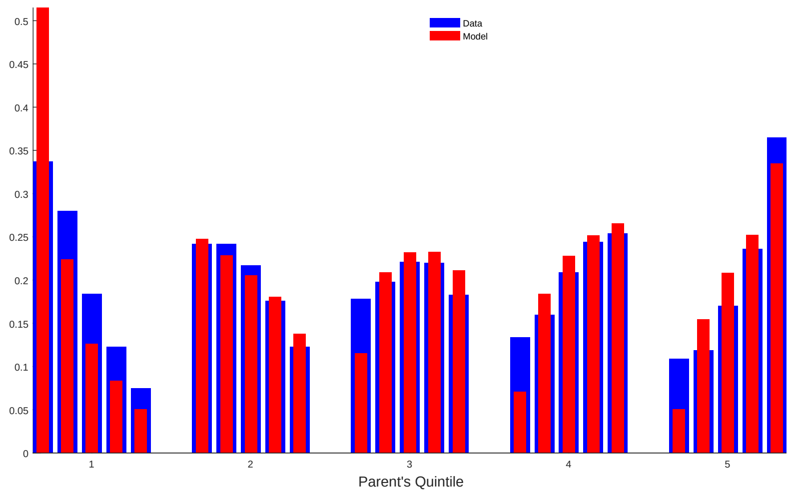

Finally, we do a very simple exercise in order to show that the model is consistent with the cross-country patterns described in Figures 1 and 3. Recall that in the benchmark calibration the wage was normalized to one. Consequently, in order to generate economies with different levels of GDP per capita as in the data, we move wages (keeping relative prices of education and childcare stable), such that the real wage (i.e. in consumption terms) is the main change. Table 6 shows that the cross-country correlation between GDP per capita and fertility elasticity and between the Gini coefficient and intergenerational mobility, both in the data and the model. This rather simple exercise shows that the model is at least capable of qualitatively capturing these correlations. Moreover, it also does a decent quantitative job, particularly with regards to the correlation between GDP per capita and fertility elasticity. We take this as evidence that the model can also capture our main patterns of interest outside of the economy on which the benchmark is estimated.

Table 6: Cross-country patterns

	Data	Model
corr(GDP pc,Fertility Elasticity)	0.773	0.827
corr(Gini,Intergenerational Mobility)	0.761	0.519

5 Results

We now use the model to answer the questions presented in the introduction. We first address how much of the inequality observed in the data is due to initial opportunities (as of age 13). Table 7 summarizes our main results. Its first four rows recall that, as discussed in the previous section, the model generates levels of inequality and social mobility consistent with the data. There are different ways to analyze the impact of initial conditions, but we focus on two that are popular in the literature.

Table 7: Social mobility and income inequality

	Data	Model
Intergenerational Mobility: Rank-Rank	0.341	0.440
Intergenerational Mobility: Log-Log	0.344	0.334
Transition Par Q1-Child Q5	0.075	0.051
Theil-L index (25-29 yo, Total)	0.202	0.154
Theil-L index relative (Parents inc)	16-75%	18.9%
Theil-L index relative (Initial Cond)	16-75%	71.6%
Variance years of education		7.236
% expl. by $h0$		77.814
% expl. by transfers		21.436

We first follow the inequality of opportunity literature (Brunori et al., 2013; Niehues and Peichl, 2014). Partitioning the population into a set of disjunct types (i.e. subgroups of the population that are homogeneous in terms of their initial conditions), the income distribution within a type is a representation of the opportunity set which can be achieved for individuals with the same initial conditions but different labor market experience. Perfect equality of opportunity is achieved if distributions are identical across types. As in this literature, we focus on the first moment of this distribution, the mean income. Measuring inequality of opportunity thus means capturing the extent to which the mean income differs across types. For this we use the Theil-L index since this can be decomposed into inequality due to initial conditions and inequality due to labor income shocks. Then, the measure of inequality of opportunity is the relative share of total inequality that can be attributed to initial opportunities. Consequently, 0% implies that all the inequality is due to adult income risk whereas 100% implies that all inequality is due to initial conditions. The empirical problem with this measure is that it is very difficult to observe initial conditions. For instance, Brunori et al. (2013) estimate a lower bound of inequality of opportunity of 16% whereas Niehues and Peichl (2014) estimate an upper bound of 75%. As explained by Brunori et al. (2013), this literature focuses on observable initial differences like gender, race or parents' income. Since the first two do not exist in our model, we replicate their exercise using the only variable in our model that could potentially be observable by them. Identifying parents' income as initial conditions, the fifth row of Table 7 shows that our estimates are very close to the lower bound from these empirical studies. However, the advantage of having a model is that we can identify the true initial conditions in our economy: parents' transfer and the initial draw of human capital. In this case, the sixth row of Table 7 shows that our actual level of inequality of opportunity (as measured by this index) is 71%, very close to the empirical upper bound.

An alternative way to evaluate the importance of initial conditions is the one followed by Huggett et al. (2011). We decompose the variance of lifetime earnings into variation due to initial conditions versus variation due to adult shocks.[35] Huggett et al. (2011) found that 61% of variation in lifetime earnings is determined by age 23. Taking into account the role of families, we find that most of this is actually determined even before high school. Table 7 reports that in our model 53% of variation in lifetime earnings is determined by initial conditions as of age 13. Regarding policy implications, our model is consistent with the early childhood investment literature which suggests that improving children's conditions very early in their lives can have a very significant in their future outcomes (Gertler et al., 2013).

[35]Such a decomposition makes use of the fact that a random variable can be written as the sum of its conditional mean plus the variation from its conditional mean. As these two components are orthogonal, the total variance equals the sum of the variance in the conditional mean plus the variance the conditional mean.

Furthermore, we can also look at the relative importance of parents' transfers and initial ability, focusing on their effect on education choice. We decompose the variance in years of education into variation due to parents' transfers and variation due to initial ability. We find that parents' transfers play a crucial role, explaining 78% of the total variance in years of education. This is consistent with the results from Belley and Lochner (2007) who find that family income and wealth are relevant for educational choices.

Finally, we can use our model to ask how important endogenous fertility and family transfers are for social mobility. For this we keep the same model and estimated parameters from Section 4, but move away from the endogenous fertility and family transfers. First, we allow for endogenous family transfers but look at the exogenous fertility case where each parent has three children. Column 3 of Table 8 reports that in such a society without endogenous fertility social mobility would at least double for all our measures. The rank-rank and log-log intergenerational mobility measures would be halved while the probability of moving from the bottom to the top quintile would more than double. We also remark that if we were to remove old-age support - one of the main motives for fertility differentials in our model - we would arrive to a very similar conclusion. Alternatively, we can also look at the case where we allow for endogenous fertility but parents' transfers are exogenously constant.[36] Similarly to removing endogenous fertility, the fourth column shows that without endogenous transfers social mobility would also double.

To summarize, we find that initial opportunities account for more of income inequality than does adult income risk over the working life. Our results suggest that taking into account family transfers as well as the negative income-fertility relationship observed in the data is very important for the analysis of inequality and social mobility. A model interested in understanding social mobility or capturing the role of inequality of opportunities, should take fertility differentials and family transfers into account.

Table 8: Social mobility and income inequality

	Data	Benchmark	Constant Fertility $N = 3$	Constant Transfers >HS cost
Income top-bottom 25-59	5.286	5.337	4.982	5.462
Intergenerational Mobility: Rank-Rank	0.341	0.440	0.192	0.180
Intergenerational Mobility: Log-Log	0.344	0.334	0.141	0.135
Transition Par Q1-Child Q5	0.075	0.051	0.121	0.128
Var Lifetime Earnings, % expl. by initial conds		53%	53%	50%
Var Years of Education, % expl. by:				
Initial ability		20%	51%	100%
Parent transfer		78%	39%	0%

Some caveats are worth mentioning regarding our results. First, one concern left for future research is whether general equilibrium effects could alter the results. It is not evident to us that qualitative results should change significantly by endogenizing prices in our model, but it could be worth exploring. Even though this may not be theoretically difficult to solve, we have not dealt with it due to computational limitations. Second, our results on the role of initial conditions should be considered as applying as of age 13. It would be interesting to go all the way back to age 0 so as to fully understand how initial conditions are formed. Our model makes one step in that direction and takes into account how family transfers are determined. Although our model is calibrated to match the data it is still silent about the exact way the initial level of human capital is formed. Further research is needed to understand this.

[36] We can fix the amount transfered to different levels. Forcing zero transfers is not useful because in such a case no one would be able to study, which would move us to a completely different economy. For this reason, we choose a level such that children are able to go to high school, and those with high initial ability are able to go to college. As this is a partial equilibrium exercise, this makes sure that prices and other parameters still remain sensible.

6 Conclusion

This paper analyzes the roots of social immobility and income inequality, trying to disentangle the importance of differences in opportunities determined early in life relative to differences in effort and luck experienced over the working lifetime. In order to overcome data limitations, this paper uses a standard heterogeneous agent life cycle model with education and idiosyncratic risk extended to account for the role of families (through endogenous fertility and family transfers) in determining initial opportunities. The model also allows for human capital transmission from parents to children, calibrated to match that in the data. We propose that fertility differentials between rich and poor households can lead to substantial differences in the resources available for children, which can be important for their adult outcomes. Importantly, the model is able to capture cross-country evidence on the relation between fertility differentials, income inequality and social mobility. Income risk is calibrated to include total earnings variation, encompassing what may be considered both wage shocks as well as hours worked (or effort) differences. Typical statistics on adult income risk are well captured by the model, which is required for an impartial comparison of the importance of adult risk relative to initial conditions.

There are different ways to analyze the impact of initial conditions but we broadly find that initial opportunities (as of age 13) account for more of income inequality than does adult income risk over the working life. If we particularly look at inequality of opportunities as computed by Brunori et al. (2013) or Niehues and Peichl (2014), our model indicates that 71% of observed inequality is due to differential initial opportunities (very close to their empirically estimated upper bound). On the other hand, if we focus on the variation of lifetime earnings, as does Huggett et al. (2011), our model suggests that 53% of this is due to differences in conditions determined even before high school. Our model also points that fertility differentials and family transfers generate over 50% of the social immobility observed in the data, implying that those interested in understanding social mobility may need to take these two forces into account.

Some comments on policy implications are appropriate now. Our results suggest that improving access to education (through education subsidies for example) may have an important effect in reducing the importance of family transfers, helping reduce income inequality and improve social mobility. Alternatively, old-age support is a motive in our model for fertility differentials which introduces a new potential gain from a generous government support for the elderly. Standard analysis of retirement benefits abstracts from the long run effects these may have on inequality and social mobility through fertility and education. Our model suggests that increasing the retirement benefits (or their progressivity) could reduce the incentives of poor families to have many children. By reducing the share of children born with few resources (and who cannot access education), it may decrease inequality or improve social mobility. Finally, even though our model is silent about the forces determining the initial level of human capital (though calibrated such that the correlation between parents' and children's human capital holds as in the data), it can still shed light on its importance for the levels of inequality observed in the data. Research on the determinants of initial human capital (Cunha et al., 2010) or on ways to improve it (like Gertler et al. (2013) on early childhood investment) is needed to understand how this initial distribution may be modified to reduce later income inequality or promote social mobility.

References

B. Abbott, G. Gallipoli, C. Meghir, and G. L. Violante. Education Policy and Intergenerational Transfers in Equilibrium. NBER Working Papers 18782, National Bureau of Economic Research, Inc, Feb. 2013.

M. A. Aguiar and M. Bils. Has Consumption Inequality Mirrored Income Inequality? NBER Working Papers 16807, National Bureau of Economic Research, Inc, Feb. 2011.

A. Alesina and P. Giuliano. *Preferences for Redistribution*, pages 93–132. North Holland, 2011.

D. Altig and S. J. Davis. Borrowing constraints and two-sided altruism with an application to social security. *Journal of Economic Dynamics and Control*, 17(3):467–494, May 1993.

J. D. Angrist and W. N. Evans. Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size. *American Economic Review*, 88(3):450–77, June 1998.

R. J. Arneson. Equality and equal opportunity for welfare. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 56(1):pp. 77–93, 1989. ISSN 00318116.

A. B. Atkinson, T. Piketty, and E. Saez. Top Incomes in the Long Run of History. *Journal of Economic Literature*, 49(1):3–71, March 2011.

D. H. Autor, L. F. Katz, and M. S. Kearney. Trends in U.S. Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics*, 90(2):300–323, May 2008.

A. Banerjee, X. Meng, T. Porzio, and N. Qian. Aggregate Fertility and Household Savings: A General Equilibrium Analysis using Micro Data. NBER Working Papers 20050, National Bureau of Economic Research, Inc, Apr. 2014.

R. J. Barro and G. S. Becker. Fertility Choice in a Model of Economic Growth. *Econometrica*, 57(2):481–501, March 1989.

P. Belley and L. Lochner. The Changing Role of Family Income and Ability in Determining Educational Achievement. *Journal of Human Capital*, 1(1):37–89, 2007. URL http://ideas.repec.org/a/ucp/jhucap/v1i1y2007p37-89.html.

M. Boldrin, M. D. Nardi, and L. E. Jones. Fertility and Social Security. NBER Working Papers 11146, National Bureau of Economic Research, Inc, Feb. 2005.

P. Brunori, F. H. G. Ferreira, and V. Peragine. Inequality of Opportunity, Income Inequality and Economic Mobility: Some International Comparisons. Working Papers 284, ECINEQ, Society for the Study of Economic Inequality, Jan. 2013.

R. Chetty, N. Hendren, P. Kline, and E. Saez. Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. NBER Working Papers 19843, National Bureau of Economic Research, Inc, Jan. 2014.

G. A. Cohen. On the currency of egalitarian justice. *Ethics*, 99(4):pp. 906–944, 1989. ISSN 00141704.

M. Corak. Income Inequality, Equality of Opportunity, and Intergenerational Mobility. *Journal of Economic Perspectives*, 27(3):79–102, Summer 2013.

D. Cox and E. Jimenez. Achieving Social Objectives through Private Transfers: A Review. *World Bank Research Observer*, 5(2):205–18, July 1990.

F. Cunha, J. J. Heckman, and S. M. Schennach. Estimating the Technology of Cognitive and Noncognitive Skill Formation. *Econometrica*, 78(3):883–931, 05 2010.

D. de la Croix and M. Doepke. Inequality and Growth: Why Differential Fertility Matters. *American Economic Review*, 93(4):1091–1113, September 2003.

G. Fella. A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344, April 2014.

R. Fernndez and R. Rogerson. Sorting And Long-Run Inequality. *The Quarterly Journal of Economics*, 116(4): 1305–1341, November 2001.

R. Fernndez, N. Guner, and J. Knowles. Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality. *The Quarterly Journal of Economics*, 120(1):273–344, January 2005.

P. Gertler, J. Heckman, R. Pinto, A. Zanolini, C. Vermeersch, S. Walker, S. M. Chang, and S. Grantham-McGregor. Labor Market Returns to Early Childhood Stimulation: a 20-year Followup to an Experimental Intervention in Jamaica. NBER Working Papers 19185, National Bureau of Economic Research, Inc, June 2013.

J. J. Heckman, L. J. Lochner, and P. E. Todd. *Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond.* , volume 1 of *Handbook of the Economics of Education*, chapter 7, pages 307–458. Elsevier, June 2006.

M. Huggett, G. Ventura, and A. Yaron. Sources of Lifetime Inequality. *American Economic Review*, 101(7): 2923–54, December 2011.

C. I. Jones and J. Kim. A Schumpeterian Model of Top Income Inequality. NBER Working Papers 20637, National Bureau of Economic Research, Inc, Oct. 2014.

L. E. Jones and A. Schoonbroodt. Complements versus Substitutes and Trends in Fertility Choice in Dynastic Models. Technical report, 2009.

L. E. Jones and M. Tertilt. *An Economic History of Fertility in the U.S.: 1826-1960.* , chapter 5, pages 165–230. Frontiers of Family Economics. Emerald, 2008.

L. E. Jones, A. Schoonbroodt, and M. Tertilt. Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? In *Demography and the Economy*, NBER Chapters, pages 43–100. National Bureau of Economic Research, Inc, May 2010.

M. Kremer and D. L. Chen. Income Distribution Dynamics with Endogenous Fertility. *Journal of Economic Growth*, 7(3):227–58, September 2002.

D. Krueger and A. Ludwig. Optimal Capital and Progressive Labor Income Taxation with Endogenous Schooling Decisions and Intergenerational Transfers. Working papers, 2015.

C. Lenihan. *IPEDS Analytics: Delta Cost Project Database 1987-2010.* National Center for Education Statistics, Institute of Education Sciences, U.S. Department of Education. Washington, DC, 2012.

S. Lundberg and R. A. Pollak. The American Family and Family Economics. *Journal of Economic Perspectives*, 21(2):3–26, Spring 2007.

R. E. Manuelli and A. Seshadri. Explaining International Fertility Differences. *The Quarterly Journal of Economics*, 124(2):771–807, May 2009.

E. G. Mendoza, A. Razin, and L. L. Tesar. Effective tax rates in macroeconomics: Cross-country estimates of tax rates on factor incomes and consumption. *Journal of Monetary Economics*, 34(3):297–323, December 1994.

B. D. Meyer and J. X. Sullivan. Consumption and Income Inequality and the Great Recession. *American Economic Review*, 103(3):178–83, May 2013.

R. Michael and C. Citro. *Measuring Poverty: A New Approach.* National Academies Press, 1995. ISBN 9780309051286.

D. Mookherjee, S. Prina, and D. Ray. A Theory of Occupational Choice with Endogenous Fertility. *American Economic Journal: Microeconomics*, 4(4):1–34, November 2012.

O. F. Morand. Endogenous Fertility, Income Distribution, and Growth. *Journal of Economic Growth*, 4(3): 331–49, September 1999.

J. Niehues and A. Peichl. Upper bounds of inequality of opportunity: theory and evidence for Germany and the US. *Social Choice and Welfare*, 43(1):73–99, June 2014.

J. B. Nugent. The Old-Age Security Motive for Fertility. *Population and Development Review*, 11:75–97, March 1985.

V. Quadrini and J.-V. Rios-Rull. *Inequality in Macroeconomics.* . Handbook of Income Distribution. North-Holland, 2014.

N. Roys and A. Seshadri. On the Origin and Causes of Economic Growth. 2014.

A Additional Figures and Tables

Table A.1: Countries included in empirical work

Country	Years	Source
Austria	1987	LIS
Brazil	1970, 1980, 1991, 2000, 2010	IPUMS
Canada	1991, 1994, 2000, 2004, 2007, 2010	LIS
China	2002	LIS
Colombia	2004, 2007, 2010	LIS
Czech Republic	1992	LIS
Denmark	1987, 1992, 1995, 2000, 2004, 2007, 2010	LIS
Dominican Republic	1981, 2002	IPUMS
Finland	1987, 1991, 1995, 2000, 2004, 2007	LIS
France	1978, 1984, 1989, 1994, 2000	LIS
Germany	1973, 1978, 1983, 2000, 2004	LIS
Guatemala	2006	LIS
India	2004	LIS
Italy	1986, 1987, 1989, 1991, 1993, 1995, 1998, 2000	LIS
Mexico	1984, 1989, 1992, 1994, 1996, 1998, 2002, 2004, 2008, 2010	LIS
Mexico	1995, 2000	IPUMS
Panama	1980, 1990, 2010	IPUMS
Peru	2004	LIS
Poland	1995, 1999, 2004, 2007, 2010	LIS
Puerto Rico	1970, 1980, 1990, 2000, 2005	IPUMS
Romania	1995, 1997	LIS
Slovak Republic	1992	LIS
South Africa	1996, 2001, 2007	IPUMS
South Korea	2006	LIS
Spain	1980, 1990, 2004, 2007, 2010	LIS
Taiwan	1981, 1986, 1991, 1995, 1997, 2000, 2005, 2007, 2010	LIS
United States	1960, 1970, 1980, 1990, 2000, 2005, 2010	IPUMS
United Kingdom	1994, 1999, 2004, 2007, 2010	LIS
Uruguay	2004	LIS

Sources: Minnesota Population Center (IPUMS) and Luxembourg Income Study Database (LIS)

Figure A.1: Income Elasticity of Fertility and GDP: All observations

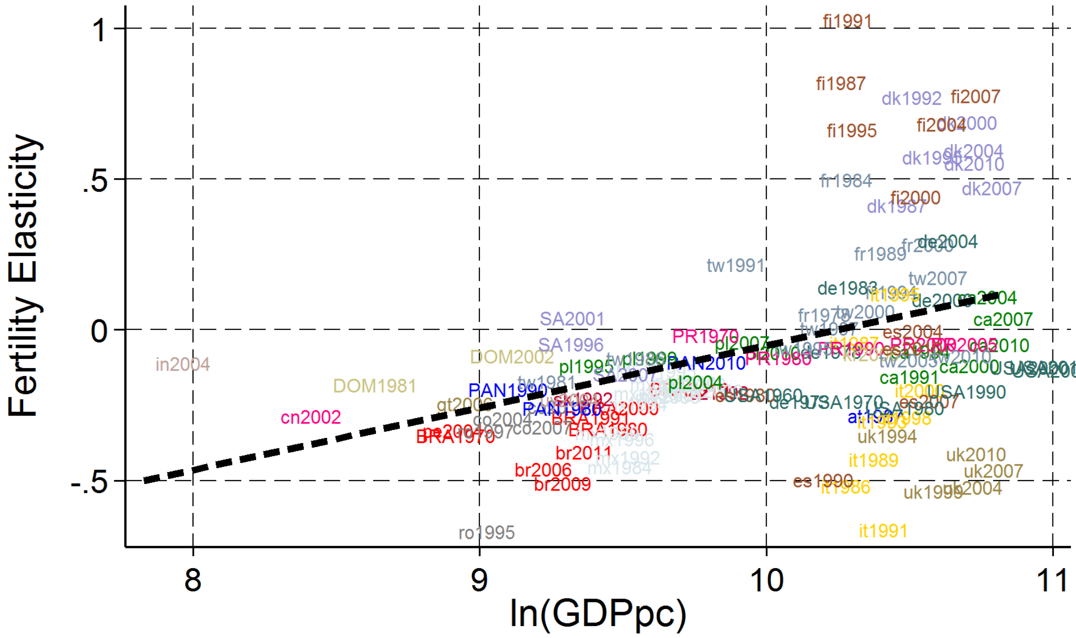

Source: IPUMS International. Methodology is explained in the main text.

Figure A.2: Income Elasticity of Fertility and Inequality: All observations

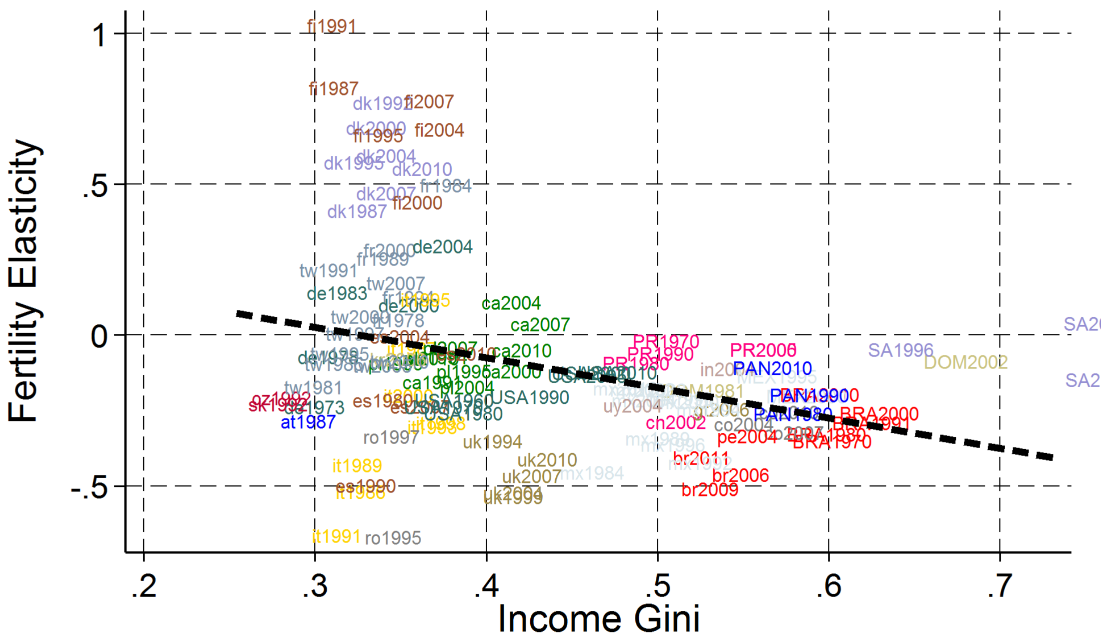

Source: IPUMS International. Methodology is explained in the main text.

Figure 1 showed that income inequality differs across countries and that it seems associated with lower social mobility. In the current Section we have shown that poorer countries tend to have bigger fertility differentials than richer ones. How does income inequality relate to fertility differentials? Intuitively, we would expect that in countries with bigger fertility differentials, stronger inequalities are present from the earliest stages of life. As poor families have more children and fewer resources to split than richer ones, many children are born with very scarce resources while a smaller group is born with a bigger pie to split among fewer hands. Assuming this affects their education or business opportunities, we would expect to observe more inequality in countries with bigger fertility differentials. Figure A.3 looks at these using the same sample as before and calculating income inequality using total household income. Then we divide the observations into deciles according to their levels of income inequality. For each decile we calculate the median level of inequality as well as the median fertility elasticity of income. As expected, Figure A.3 shows that countries with more inequality are associated with higher fertility differentials than more equal ones. This evidence coincides with that shown by Kremer and Chen (2002) which looked at fertility differentials across education groups.

Figure A.3: Income Elasticity of Fertility by Income Gini Decile

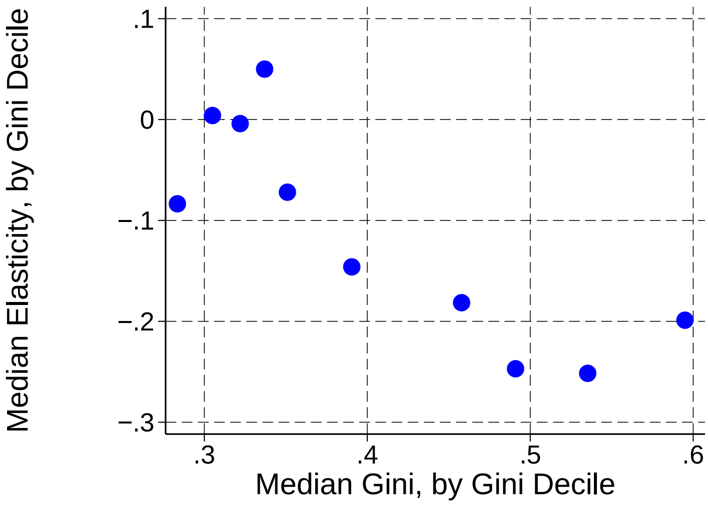

Source: IPUMS International and LIS. Methodology is explained in the main text. Figure A.2 in the Appendix includes all the data observations before grouping them in deciles.