Macroeconomic Effects of 'Free' Secondary Schooling in the Developing World

Junichi Fujimoto (GRIPS)
David Lagakos (BU and NBER)
Mitchell VanVuren (Yale)

March 2024

Abstract

This paper studies the macroeconomic effects of publicly funded ('free') secondary schooling in the developing world. Our analysis is based on an overlapping generations model of human capital accumulation that we estimate to match experimental evidence on the effects of scholarships for poor but talented students in Ghana. The model predicts that nationwide free secondary schooling increases average education levels but reduces GDP per capita in the long run. The human capital gains from free schooling in the model are offset by lost income during schooling years, negative selection of new students, and reductions in fertility by high-ability households.

Email: j-fujimoto@grips.ac.jp, lagakos@bu.edu, mitchell.vanvuren@yale.edu. For helpful comments we thank Ufuk Akcigit, Marco Bassetto, Natalie Bau, Kevin Donovan, Raquel Fernández, Chang-Tai Hsieh, Aubhik Khan, Jeremy Pearce, Karthik Muralidharan, Michael Peters, Tommaso Porzio, Pascual Restrepo, Todd Schoellman, Julia Thomas and seminar/conference participants at Berkeley, Chicago, City University of Hong Kong, the Development Bank of Japan, the Federal Reserve Banks of Minneapolis and St. Louis, Georgetown, Hitotsubashi, Kyoto, NYU, Ohio State, Penn State, Riverside, Tohoku, the Cowles Foundation Conference on Structural Change, the UCLA Human Capital and Development Conference and the NBER Summer Institute. We thank the IGC, JSPS (Kakenhi No. 18K01505 and No. 22K01399), and GRIPS Policy Research Center for financial support and Ariel Akerman, Soala Ekine, and Isaac Otoo for outstanding research assistance. All potential errors are our own.

1. Introduction

It has been said that talent is universal but opportunity is not. This saying could hardly apply better than to the millions of school-aged children throughout the developing world who are not enrolled in school. A common explanation of the low secondary schooling levels is credit constraints that prevent poorer households from borrowing for education. Keeping bright young people out of secondary school may lead to a significant misallocation of talent in the education system, which can reduce aggregate productivity levels (Hsieh, Hurst, Jones, and Klenow, 2019). More generally, low levels of human capital in developing countries are thought to be one of the most important proximate causes of their low income levels (e.g. Hall and Jones, 1999; Bils and Klenow, 2000; Erosa, Koreshkova, and Restuccia, 2010; Hanushek and Woessmann, 2012; Manuelli and Seshadri, 2014; Hendricks and Schoellman, 2018) as well as a barrier to structural change (Porzio, Rossi, and Santangelo, 2022; Buera, Kaboski, Rogerson, and Vizcaino, 2022).

In this paper, we analyze the macroeconomic effects of an increasingly common policy aimed at raising secondary enrollments: publicly funded ('free') secondary schooling. A number of low-income countries have recently adopted free schooling in the hopes of raising human capital levels, and others are seriously considering it (see Center for Global Development, 2022). To our knowledge, however, ours is the first macroeconomic study of free secondary schooling in the developing world.

Our analysis is based on an overlapping generations (OLG) model of human capital accumulation with heterogeneous households and credit constraints in the spirit of Galor and Zeira (1993) and Bénabou (2002). Our model incorporates three key features necessary for understanding the long-run general-equilibrium impacts of education policy in the developing world. The first is opportunity cost of schooling in the form of lost labor income for those attending secondary schooling rather than working. It is hard to ignore this opportunity cost since people of secondary-school age are close to their prime working years in developing countries. The second feature is a secondary-school entrance exam. This merit requirement is relevant since secondary school currently selects students with relatively high ability, meaning that free schooling for all could lower average student ability levels. The third feature is fertility that depend on education, and in particular lower fertility for those with more education. Evidence from a large set of developing economies has linked education expansions and economic growth to significant declines in fertility (Chatterjee and Vogl, 2018).

We estimate the model using experimental evidence from a randomized controlled trial (RCT) that offered secondary school scholarships to a randomly selected set of poor but high-ability students in Ghana (Duflo, Dupas, and Kremer, 2021). Those offered the scholarships were about 25 percent more likely to finish secondary school than a control group four years hence. Scholarship winners performed about 0.2 standard deviations higher on tests of literacy and mathematics, which is comparable to the effects found in other successful education interventions. Moreover, scholarship winners had significantly fewer children than the control group.

We use our estimated model to simulate the aggregate effects of free secondary schooling in general equilibrium. The model predicts a one-third increase in the number of secondary school graduates, which is broadly consistent with the substantial increases in attendance described by the Center for Global Development (2022). Yet the model predicts a *negative* effect on GDP per capita, which falls by around 1 percent in the long run. Our model's confidence interval for this prediction is entirely negative, suggesting that even modest gains in GDP per capita from free secondary schooling are unlikely. Taxes increase by about 1 percent per capita in the long-run, meaning that the policy pays for only a fraction of its costs.

To better understand why the model makes such pessimistic quantitative predictions we conduct a series of counterfactual experiments. In the first, we eliminate the opportunity cost of secondary schooling by exogenously disallowing work by those of secondary school age. We then simulate the impacts of free secondary schooling on this economy. We find that it would have raised GDP per capita by about 3 percent, signaling that the earnings gains from free school are largely offset by lost earnings during secondary school years in our main analysis.

In a second counterfactual, we simulate a hypothetical free schooling policy that maintains a similar skill distribution as the poor but talented students studied in the RCT. Free schooling causes no change in GDP per capita in this economy, pointing to the less favorable selection of students once everyone is allowed to attend secondary school (like Hendricks and Schoellman (2014) find for college in the United States).

In a third counterfactual we eliminate the differential fertility rate by education level, meaning there is no way the population share of the educated can fall after education levels rise. In this counterfactual, GDP per capita rises by about 1 percent in the long run due to free secondary school, highlighting how long-run changes in fertility in the main analysis play a quantitatively important role in lowering GDP per capita once secondary schooling is made free.

When we simulate the counterfactual impacts of shutting down all three channels we find that free secondary schooling would lead to a GDP per capita gain of around 7 percent in the long run. Taxes per capita would change little in this case as free schooling pays for itself in the long run. The sizable GDP per capita gain in this counterfactual illustrates how the model's pessimistic conclusions about free schooling are not hard wired into the findings of the scholarships RCT, but are rather due to the economic forces of opportunity cost, negative selection of new graduates, and reductions in fertility by high-ability households.

An alternative class of education policies studied in the development literature have focused on raising the quality of schooling, rather than making schools more accessible. We use our model to simulate an economy-wide improvement in schooling quality, which could represent pay-for-performance incentives for teachers (Muralidharan and Sundararaman, 2011; Duflo, Hanna, and Ryan, 2012; Mbiti, Muralidharan, Romero, Schipper, Manda, and Rajani, 2019), additional teachers in the classroom (Banerjee, Cole, Duflo, and Linden, 2007) or other interventions shown to bolster student academic performance. We find that school quality improvements are significantly more effective at raising average income levels than free secondary schooling. A nationwide school quality improvement raising test scores by 0.1 standard deviations – consistent with the experiments above – leads to a long-run GDP increase of 3 percent per capita, and pays for itself in equilibrium. We conclude that improving the quality of secondary schools in Africa is a more effective way to raise living standards than making the current schools free.

A welfare analysis of our model predicts that the richest quartile of households suffer welfare losses from free schooling, largely since they pay higher taxes, whereas the poorest quartile gain relatively most. We conclude that free secondary schooling is predominantly redistributionary in nature, and that low secondary schooling levels are largely an efficient response to low quality school options. This implication is broadly in line with the conclusions of the macro development literature emphasizing low schooling quality, rather than low average years of schooling, as the proximate cause of low human capital levels in poor countries (Hanushek and Woessmann, 2007; Schoellman, 2012).

Related Literature. This paper belongs to the vast literature on macroeconomic models of human capital investment (e.g. Becker and Tomes, 1979; Loury, 1981; Galor and Zeira, 1993; Restuccia and Urrutia, 2004; Erosa et al., 2010; Manuelli and Seshadri, 2014; Lee and Seshadri, 2019). In particular, our quantitative exercises build on the literature on the macroeconomic effects of credit constraints in education in advanced countries, such as the seminal work of Lochner and Monge-Naranjo (2011). Our paper is most closely related to the studies by Abbott, Gallipoli, Meghir, and Violante (2019) and Daruich (2020), both of whom study expansions in publicly funded education in the United States. Both studies reach fairly positive conclusions about the effects of expanding public education, unlike our study, which arguably reflects disparities in school quality between rich and poor countries. As in Daruich (2020), we discipline our model using experimental evidence from an RCT.

Similar to our paper, Khanna (2023) studies the impact of expanding education in general equilibrium and finds a large role for relative wage effects. Comparatively, we find a smaller role for wage impacts and instead emphasize the importance of other channels such as fertility and opportunity cost.

Our paper also builds on the recent literature attempting to quantify the extent to which credit market imperfections drive misallocation in developing countries. Bassi, Muoio, Porzio, Sen, and Tugume (2022) and Caunedo and Kala (2022) show how rental markets for large indivisible capital goods can reduce capital misallocation, and Moll (2014) and Midrigan and Xu (2014) find a significant ability for firms to save their way around credit constraints.

In estimating our model to a field experiment, we build on a growing body of macroeconomic research on development that uses randomized experiments in order to guide general-equilibrium counterfactuals (see Buera, Kaboski, and Townsend, 2023). Our paper is the first to take this approach when studying the macroeconomic effects of education policy in the developing world. Other studies using this methodology have studied small business investment (Kaboski and Townsend, 2011), occupational choice (Buera, Kaboski, and Shin, 2021), infrastructure investments (Brooks and Donovan, 2020), rural-urban migration (Lagakos, Mobarak, and Waugh, 2023), and firm training programs (Akcigit, Alp, and Peters, 2021). Our methodology relates to the papers by Todd and Wolpin (2006) and Attanasio, Meghir, and Santiago (2012) that use structural labor models to interpret experimental evidence; see Todd and Wolpin (2023) for an engaging recent review.

2. Secondary Schooling in the Developing World

In this section we summarize the macro and micro evidence on secondary schooling outcomes and free secondary schooling policies in the developing world. We also discuss how this evidence informs our modeling choices and counterfactual simulations in the sections that follow.

Figure 1: Primary and Secondary School Enrollment Rates

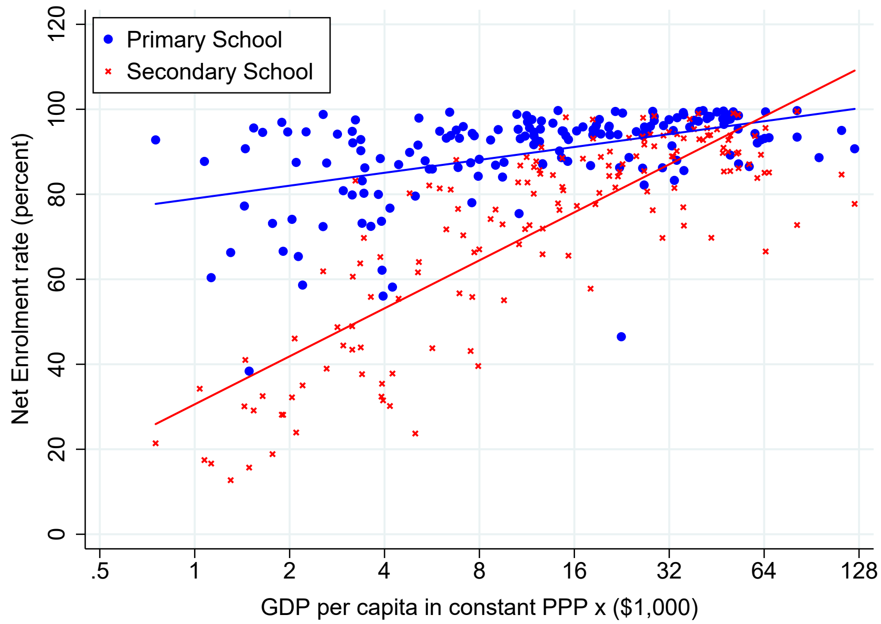

Aggregate data on schooling enrollment show plainly that developing countries mainly lag behind richer ones when it comes to secondary schooling (as opposed to primary schooling). Figure 1 plots net enrollment rates in primary school (blue dots) and secondary school (red x's) in 2019 against GDP per capita using data from the World Bank. Net enrollment rates are defined as the number of people enrolled in school relative to the population of school-aged individuals. In the world's poorest countries, roughly four out of five children of primary-school age are enrolled in school, compared to nearly every child in the richest ones. For secondary schooling, the differences are much starker. At the bottom of the world income distribution, only around one-third of those of secondary-school age are enrolled in secondary school, whereas at the top, enrollment rates are again near one hundred percent.

One salient difference between rich and poor countries in terms of education policy is that richer countries are much more likely to publicly finance secondary education. It is not surprising then that many developing countries have recently considered implementing 'free' schooling policies, in which the government finances school fees for at least some secondary-age students (see Center for Global Development (2022) and Appendix Table A.1, which lists the developing countries, like Ghana, that already have some form of free secondary schooling in place). One main rationale for publicly funded schooling is to help raise average schooling levels and hence GDP per capita. A second rationale is to make secondary education more accessible to poorer households, consistent with redistributionary motives. These two objectives are not necessarily in contrast with one another, since raising average years of schooling is likely to require expanding schooling access to poorer households that were previously unable to pay for secondary school fees.

Recently, a number of micro studies have estimated the impacts of merit-based scholarship programs in developing countries, though with mixed results. Brudevold-Newman (2021) found, using a difference-in-difference approach, that free secondary schooling in Kenya increased educational attainment, reduced fertility, and increased the likelihood of skilled work. Using a regression discontinuity design in Cambodia, Filmer and Schady (2014) found that scholarships increased educational attainment but did not increase earnings, fertility, or test scores. Both studies highlight credit constraints as a reason more students were not already enrolled.

Duflo, Dupas, and Kremer (2021) conducted the first long-run RCT evaluation of a merit-based scholarship program for secondary school. Their study is set in Ghana, where the education system consists of primary school and junior high school (JHS) until age 14, at which point students are required to pass the Basic Education Certificate Examination (BECE) in order to attend senior high school (SHS). The authors identified approximately two thousand students who had passed the BECE in 2008 but had not enrolled in SHS by the deadline for the next school year. Among these students, one-third were randomly selected to receive a four-year scholarship covering one hundred percent of tuition and fees.

Students who received a scholarship were substantially more likely to complete SHS relative to those in the control group, and on average completed 1.2 more years of school than the control group. Scholarship winners similarly exhibited higher human capital, as measured by math and reading tests. The magnitudes were substantial, and translate into 7.6 percent higher human capital per year of school completed. There was also a significant negative impact on fertility, with scholarship winners having fertility declines of around 10 percent after 12 years. We view this experiment as the most comprehensive and credible evaluation of free secondary schooling to date. Consequently, we use these experimental moments to parameterize our model, which we develop in the next section.

Taking these results at face value, what might one infer about the aggregate gains from a national free schooling policy? Suppose, like in the RCT, that the policy could raise average schooling by 1.2 years among the 70 percent of Ghanaians of SHS age who do not attend school. With a human capital increase of 7.6 percent per year, consistent with impacts on test scores in the experiment described above, this would result in a roughly 6 percent increase in GDP in the long run (=0.7 x $(1.076^{1.2} - 1)$). This paints a promising picture for free secondary school, and suggests that the aggregate effects of the policy could in principle be significant.

3. Overlapping Generations Model

We now describe the model, which, at its heart, is an OLG model of human capital accumulation with credit constraints a la Galor and Zeira (1993) and Bénabou (2002). The model is tailored to capture key features of developing countries relevant for how free secondary schooling policies impact average income per capita. The model features opportunity cost of schooling, in the form of lost labor income; merit requirements to attend secondary school, modeled as a test score cutoff; and fertility levels that depend on education levels. As in most models in this literature, ours allows for misallocation of talent through borrowing constraints that keep high ability children of poor parents out of school.

We also include several other features that are relevant for estimating and quantifying the effects of free secondary schooling. Consistent with the evidence of Khanna (2023), we model the labor of different education types as imperfect substitutes, so that an increase in the supply of educated workers depresses their relative wages. As in a growing literature in macro development we allow for saving constraints, which help match the low average levels of liquid asset holdings, and impede households from simply saving themselves out of borrowing constraints (see e.g. Donovan, 2021). Following the literature on public finance and development we posit a tax system in which a narrow base of high earners pay the majority of the taxes used to finance public expenditures (see e.g. Jensen, 2022).

3.1. Environment

Time is discrete and goes from 0 to infinity. There is a single good which can be used for consumption, savings, and investment in education. The economy is populated by overlapping generations of households that are heterogeneous in their parental human capital, child ability, taste for schooling, and savings. The timeline of events for these households is shown in the graphic below.

Individuals live for 14 periods, where each period corresponds to 5 years. For their first five periods of life (ages 0–24) children live with their parents. In the third period (ages 10–14), all children attend JHS. We abstract from the choice of attending school at this age based on the evidence of the previous section that virtually all children already attend JHS. In the fourth period, (ages 15–19), children either attend SHS (i.e. secondary school) or work. This is the key schooling choice in the model. The fact that a household must give up a period of a child's work, and thus income, in order to attend secondary school captures the notion that there remains an opportunity cost even in the case where schooling is made free. In the fifth period, all children work using their respective education level, which is fixed for the remainder of their life.

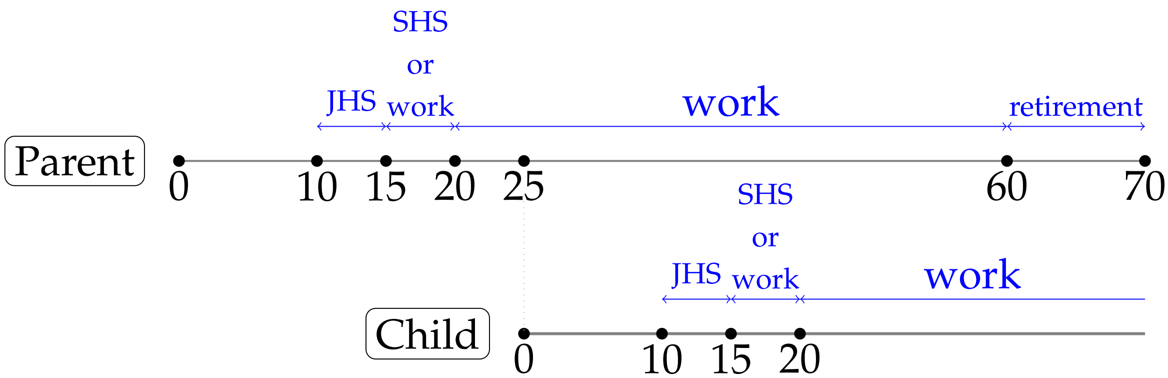

At the beginning of period six, when turning age 25, children leave their parents, have children, and become parents themselves. We abstract from household formation decisions since they do not seem crucial for our task at hand. Instead, we model households as continuous dynasties that do not mix. These new parents then work from age 25 to 60, at which point they retire, and die at age 70. This is roughly the average life expectancy in Ghana, for example, whose features we will use to parameterize our model in the following section.

Each new household consists of a parent aged 25 and newborn children. The model features population growth, and the number of children, denoted as $1 + \nu_{s_p}$, decreases with the parent's schooling level s_p. As a result, policies that increase an individual's level of schooling will decrease their fertility, consistent with a variety of evidence for developing countries.

Individuals are heterogeneous in learning ability $z \in Z = \{z^1, z^2, \ldots, z^N\}$. The ability within a household follows a first-order Markov chain which, by Tauchen's (1986) method, approximates the AR(1) process:

$$\log z_c = \rho \log z_p + \upsilon,\ \rho \in (0, 1)\,. \tag{1}$$

Here, z_p and z_c denote the parent and children's ability. Throughout, variables with subscripts p and c pertain to parents and children, respectively. The random shock υ is independently and identically distributed (i.i.d.) and follows $N\left(0, \sigma_{\upsilon}^{2}\right)$. Thus, ability is transmitted within each household but only imperfectly, and is identical across siblings. Following the evidence in e.g. Cunha and Heckman (2007), we interpret ability to be a function of inherited capabilities and parental inputs.

All household decisions are made by parents, who derive flow utility $U\left(c\right) = \log\left(c\right)$ from household consumption $c \geq 0$ and discount the future at rate $\beta \in (0, 1)$. The assumption that parents make educational decisions follows the tradition of Becker and Tomes (1979) and is consistent with evidence that parents in low-income countries predominantly take an authoritarian approach to parenting, dictating decisions directly rather than trying to reach an agreement with children (Doepke and Zilibotti, 2017). Parents and children (from ages 15 to 25) have a single unit of time each period which they supply inelastically to wage work or education. Parents are imperfectly altruistic toward children and therefore derive utility also from children's well-being (as in e.g. Laitner, 1997).

Parents make schooling decisions for their children when the children turn age 15, after observing the children's ability and test scores as well as the children's realization of a(schooling taste shock. More precisely, children enjoy random utility (internalized by the parent through imperfect altruism) δ_s from schooling level $s \in S = \{J, S\}$ (JHS, SHS), where δ_s follows a standard Gumbel distribution with scale parameter θ. Parents must forgo a period of children's income to send their children to an additional period of school, and further, providing children final schooling level $s \in S$ requires goods costs Ψ_s. These goods costs represent school fees and satisfy $\Psi_S > \Psi_J = 0$, where the equality reflects the free primary education that prevails in most developing countries. Thus when deciding whether to send their children to school, parents must consider the opportunity cost of their children's work as well as the explicit goods cost.

Households face incomplete markets as in Aiyagari (1994), Bewley (1977), and Huggett (1993) and cannot borrow but can save at an exogenous rate r. While households do face idiosyncratic income risk, the most important feature of this borrowing constraint is that it prevents parents from borrowing against their child's future income in order to fund school attendance. This allows for the possibility that a high ability child, whose return to additional schooling far exceeds the cost, may not attend if born to a poor parent, resulting in misallocation.

To capture the fact that one must pass an entrance test to enter secondary schooling in most developing countries, we set a threshold test score for entering SHS. One's test score $\tilde{z}$ is related to ability as

$$\tilde{z} = z + \varepsilon, \tag{2}$$

where the noise ε follows $N(0, \sigma_\varepsilon^2)$.

The human capital of an individual with ability z and schooling s is given by

$$h(z, s) = \begin{cases} 1 & \text{if } s = J, \\ z \cdot \eta_S & \text{if } s = S, \end{cases} \tag{3}$$

where $\eta_S > 0$ and represents the efficiency, or quality, of schooling. Thus, ability affects human capital only for those with SHS education, and the resulting human capital of a secondary education depends on the product of the student's ability and the schooling quality.

Markets are competitive and the aggregate production function, operated by a representative profit-maximizing firm, is given by:

$$Y = AK^\alpha \left[(N_J)^\lambda + (N_S)^\lambda \right]^{\frac{1-\alpha}{\lambda}}, \ \alpha, \lambda \in (0, 1). \tag{4}$$

Here, A is aggregate productivity, K is physical capital, and N_s is aggregate efficiency units of labor of individuals with schooling level s. The firm rents physical capital from households or foreign investors at an exogenous international market rate r^*. Due to savings frictions, however, the return to physical capital for households is lower, at $r = r^* - \chi < r^*$. This lower return to capital helps us match the low savings rates among households in low-income economies (as in Donovan, 2021).

The labor income y of an individual equals the product of three terms. The first term is the wage rate per efficiency units of unskilled ($s = J$) or skilled ($s \in S$) labor, denoted as w^U or w^S, respectively. The second term, ζ, represents idiosyncratic shocks to labor productivity, where $\log \zeta$ is drawn each period from $N\left(0, \sigma_\zeta^2\right)$. The third term is human capital $h(z, s)$, given by (3). For example, the labor income of an individual with education level S is given by

$$y(z, S, \zeta) = w^S \zeta h(z, S) = w^S \zeta z \eta_S. \tag{5}$$

3.2. Parents' Problems

Parents make consumption and saving decisions in each period, and additionally, schooling decisions when their children reach the age for secondary school. We discuss below the parents' problems in the key periods in the life-cycle; we omit the description of their problems in other periods, which are standard consumption-savings problems. In addition to individual state variables described below, the parent's problems depend on the p.d.f. f describing the distribution of households across individual states and the aggregate population level P.

When $\tau = 9$, and children turn 15, parents observe the realizations of the schooling taste shocks (δ_J, δ_S), children's ability and test score $(z_c, \tilde{z}_c)$, and their own and children's labor productivity (ζ_p, ζ_c). Then, if $\tilde{z}_c$ weakly exceeds the threshold test score $\bar{z}$, parents have the option of sending children to an additional period of schooling ($s = S$). The value function of such parents with ability z_p, schooling level s_p, and assets a is given by

$$\begin{aligned} V_9(a, z_p, s_p, \zeta_p, \delta_J, \delta_S, z_c, \zeta_c; f, P|\tilde{z}_c \geq \bar{z}) = \max_{c \geq 0, a' \geq 0, s'_c \in \{J,S\}} \log(c) \\ + \delta_J \mathbf{I}\left(s'_c = J\right) + \delta_S \mathbf{I}\left(s'_c = S\right) + \beta \mathrm{E}\left[V_{10}(a', z_p, s_p, \zeta'_p, z_c, s'_c, \zeta'_c; f', P')\right] \end{aligned}$$

where the maximization is subject to the flow budget constraint

$$\begin{aligned} a' + c + \left(1 + \nu_{s_p}\right)\mathbf{I}\left(s'_c = S\right)\Psi_S \\ = \\ y_p(z_p, s_p, \zeta_p) + (1+r)a + \left(1 + \nu_{s_p}\right)\left(1 - \mathbf{I}\left(s'_c = S\right)\right)y_c(z_c, J, \zeta_c) - T(z_p, s_p, \zeta_p, z_c, J, s'_c, \zeta_c) \end{aligned} \tag{6}$$

and the perceived laws of motion for the aggregate state variables f and P, given by $f' = F(f, P)$ and $P' = H(f, P)$, respectively. Here, the prime denotes values of variables in the next period and T is total amount of taxes paid by the household, which depends on the parent and children's labor income, and is therefore a function of $(z_p, s_p, \zeta_p, z_c, J, s'_c, \zeta_c)$.[1] We suppress the dependence of y_p, y_c, and T on f and P except where it is necessary to make that dependence explicit.

When $\tau = 10$, children live one final period with their parents and work with the human capital given by their education decision the previous period. The value

[1]Note that T depends on both the children's current schooling level s_c $(= J)$ and next period's schooling level s'_c. This is because the labor income depends on educational attainment, and only the children who do not go to school ($s'_c = s_c$) earns the labor income in the current period.

function of such parents is expressed as

$$V_{10}(a, z_p, s_p, \zeta_p, z_c, s_c, \zeta_c; f, P) = \max_{c \geq 0, a' \geq 0} \log(c) + \beta \mathrm{E}\left[V_{11}(a', z_p, s_p, \zeta_p'; f', P')\right] \\ + \beta b \left(1 + \nu_{s_p}\right) \mathrm{E}\left[V_6(0, z_c, s_c', \zeta_c'; f', P')\right] \tag{7}$$

subject to

$$a' + c = y_p(z_p, s_p, \zeta_p) + (1+r)a \\ + \left(1 + \nu_{s_p}\right) y_c(z_c, s_c, \zeta_c) - T(z_p, s_p, \zeta_p, z_c, s_c, s_c', \zeta_c), \tag{8}$$

$f' = F(f, P)$, $P' = H(f, P)$, and $s_c' = s_c$. On the right-hand side of (7), V_{11} denotes the parent's value function in the following period, which no longer depends on the ability and schooling of children who become independent from parents. The last term on the right-hand side of (7) denotes utility that imperfectly altruistic parents derive from their children's well-being, where $b > 0$ is the altruism parameter and $V_6(0, z_c, s_c', \zeta_c'; f', P')$ is the value function of children who form new households with zero assets.

3.3. Government, Taxes and Equilibrium

The government collects tax revenue from households which it then spends on "public goods," and, in the policy counterfactuals, free secondary schooling. The government budget constraint in per capita terms is given by:

$$G + \xi \int \mathbf{I}\left(s_c' = S \wedge \tau = 9\right) df = \int T(z_p, s_p, \zeta_p, z_c, s_c, s_c', \zeta_c) df \tag{9}$$

where G is spending on public goods per capita and ξ is expenditure on free secondary education per capita. Since the paper is about public financing of secondary education, and not other public expenditures, we abstract from how G affects households or producers in the economy. When we simulate the effects of free public schooling, we assume that G remains unchanged, so that any schooling subsidy must be funded through per-period adjustments in the tax function T.

We focus our quantitative analysis on the balanced growth path of the economy. We relegate the full definition of recursive competitive equilibrium and the balanced growth path to Appendix B. In essence, the balanced growth path is the equilibrium in which the aggregate population level grows at a constant rate, but the relative distribution of households across individual states is constant. In this situation, household behavior does not depend on the aggregate population level. In all of our analyses, we assume the economy starts on a balanced growth path. To examine the effects of a policy change, we introduce the policy into the balanced growth path of the economy and compute transition dynamics by calculating sequences of population growth rates and prices that converge to the new balanced growth path.

3.4. Illustration: Misallocation of Talent and Impacts of Free Schooling

The key decision for a household is whether or not their children attend secondary school. The benefit of attending is higher future wages, increasing in proportion to the ability level of the child. The costs of schooling are the goods cost, Ψ_S, and the opportunity cost, represented by foregone earnings. Neither of these costs depends directly on household characteristics, but, due to borrowing constraints, the utility cost ends up being higher for households with low income and assets, who have higher marginal utility of consumption. Intuitively, misallocation arises whenever a poor family chooses not to send a high ability child to secondary school even though doing so would increase lifetime household income.

Consider two example cases of our model that vary in the goods cost of schooling, Ψ_S, the quality of secondary education, η_S, and the savings wedge, χ. The first economy, which we call the *low misallocation economy*, features a relatively low cost of schooling, relatively low schooling quality, and a small savings wedge. As we show below, borrowing constraints bind for few households in this economy, and many households choose not to send their children to secondary schooling because the returns are low. The second economy, namely the *high misallocation economy*, has a higher cost of school, higher schooling quality, and a larger savings wedge. We take other parameters to be the same across economies.[2]

Figure 2 panels (a) and (c) plot the probability that a child attends secondary school – conditional on passing the entrance exam – as a function of the child's ability and their parent's ability (a proxy for parental income and wealth). Panel (a) represents the low misallocation case, and panel (c) represents the high misallocation case. The dotted gray line labeled *unconstrained cutoff*, marks the child ability level at which the net-present-value of the additional period of schooling is exactly equal to the total cost of attendance. If households faced no borrowing constraint, this is the ability level above which all children would attend SHS, and below which

[2] In particular, the low misallocation economy features values (Ψ_S, η_S, χ) of $(5.5, 1.5, 0.1)$ while the high misallocation economy has values $(14.0, 5.0, 0.2)$. The rest of the parameter values, which are not crucial for the conclusions in this section, can be found in Appendix Table A.2.

Figure 2: Child SHS Attendance Probability

Low Misallocation

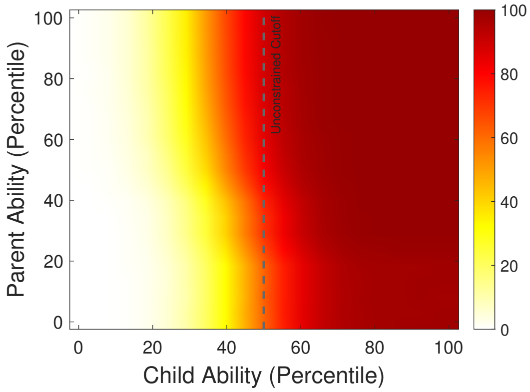

(a) SHS Attendance Probability

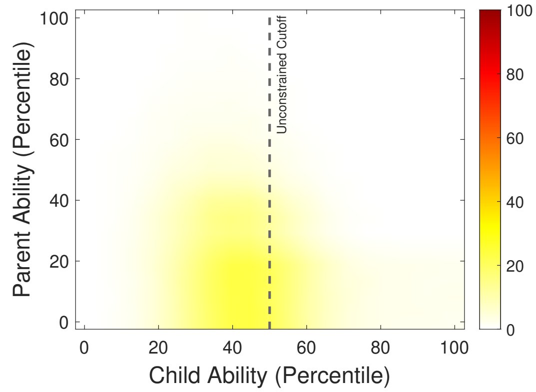

(b) Increase in Attendance from Free SHS

High Misallocation

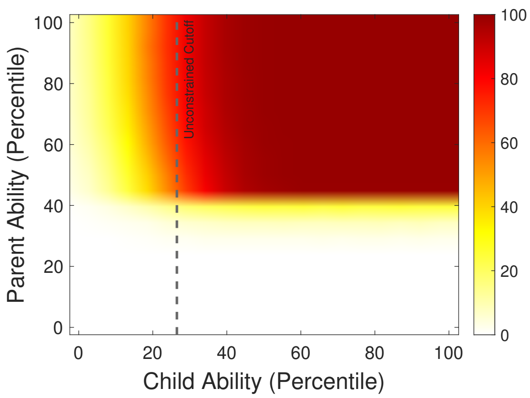

(c) SHS Attendance Probability

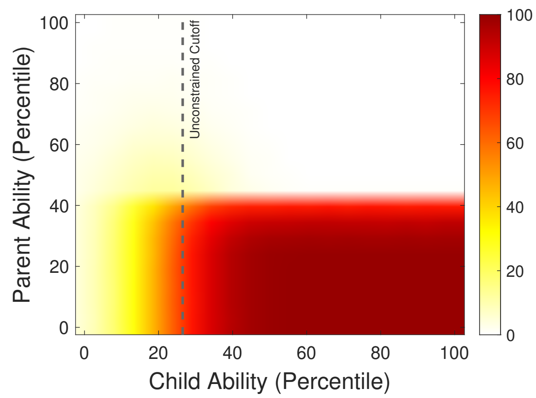

(d) Increase in Attendance from Free SHS

none would attend (assuming the average taste shock). In Panel (a), children's attendance probabilities are roughly in line with this cutoff: those above the cutoff largely attend regardless of their parent's ability, and those below largely do not.

In Panel (c), in contrast, children born to sufficiently high ability parents attend SHS roughly according to the unconstrained cutoff, and children born to low-ability parents are unlikely to attend SHS regardless of their ability level. In this economy, there is substantial misallocation in the sense that many children for whom the net-present-value of education outweighs the costs of education do not attend. As a result, one can imagine substantial scope for gains in output from reducing the cost of schooling.

Panels (b) and (d) illustrate these potential gains by displaying the increase in the probability of attendance when SHS is offered for free. In the low misallocation economy, the increase in attendance is small in magnitude and concentrated among marginal children; many of those induced to attend are actually below the unconstrained cutoff and gain little to no lifetime income from schooling. In contrast, the high misallocation economy exhibits large increases in attendance probabilities concentrated at the top of the ability distribution. These children exhibit substantial income gains from schooling, and their attendance results in higher output.

It is important to note that aggregate data alone are not necessarily informative about the reasons why secondary attendance rates are so low. Although the two economies in Figure 2 are very different, they both have aggregate attendance rates of around 30 percent. In the Panel (a) economy, attendance is low simply because the return to schooling is low on average relative to the cost, as shown by the high unconstrained cutoff. In Panel (c), however, attendance is low in large part because many high-ability children face borrowing constraints and hence forgo secondary school. This motivates our use of experimental moments, in addition to aggregate moments, in estimating the model.

4. Model Estimation

While our estimation is largely focused on experimental moments, we first choose a handful of parameters directly, either as normalizations or to match standard values from the literature. We then estimate the rest using the Simulated Method of Moments (SMM).

4.1. Directly Chosen Moments and Aggregate Moments

Table 1 presents parameters chosen directly. We start by normalizing aggregate productivity, A, to be one. We set capital's share in production, α, to be 0.33, and the discount factor, β, to be 0.96^5, which are standard values. The international market interest rate r^* is chosen to generate a (depreciation-inclusive) user cost of capital equal to 10 percent per year.

We choose the parameter governing the substitutability of skills, λ, to be 0.75, which generates an elasticity of substitution of 4. This is consistent with the long-run estimates of Bils, Kaymak, and Wu (2022) based on cross-country school attainment and wage data by attainment level. We are primarily interested in the long-run effects of schooling expansions, making a long-run elasticity of substitution appropriate for our study. We have experimented with lower values of this elasticity, down to a value of 1.4, but these do not substantively affect our main conclusions.

We pick the standard deviation of the idiosyncratic income shock, σ_ζ, to be 0.32, which matches the variance of the transitory component of log wages estimated by Lagakos and Waugh (2013). In the model, this transitory component is calculated by subtracting the variance of the permanent component of log income (explained further below) from the total variance of log income.

Finally, we set the income tax function to match Ghana's statutory income tax rates at the time of the experiment, summarized in Appendix Table A.3, which focuses on a narrow tax base consisting of only the highest income earners. This specification is consistent with the overall view that taxation in the developing world is highly progressive and absent for the poorest households (see e.g. Jensen, 2022). We then set per capita government spending on public goods G such that the government budget is exactly balanced each period along the balanced growth path.

Table 1: Directly Chosen Parameters

Description	Parameter	Value
Agg. Productivity	A	1
Share of Capital	α	0.33
Discount Factor	β	0.96^5
International Market Rate	r^*	$1.1^5 - 1$
Skill Substitutability	λ	0.75
Std. Deviation of Income Shock	σ_ζ	0.32

Note: The table reports the values of directly chosen parameters in the model.

4.2. Simulated Method of Moments

We estimate the remaining parameters of the model using the SMM. There are ten such parameters, which we estimate using ten moments. Formally, we solve for the parameter vector:

$$\Theta = \{\nu_J, \nu_S, \eta_S, \Psi_S, b, \sigma_\varepsilon, \theta, \chi, \rho, \sigma_v\} \tag{10}$$

that minimizes the sum of squared differences between the moments in Table 2 and their model counterparts. We also compute 95-percent confidence intervals for our parameters through bootstrapping, treating non-experimental moments (those above the line in Table 2) as fixed values, and re-sampling the experimental moments.

The first five moments we target do not use experimental variation; these are listed in the top portion of Table 2. The first is a population growth rate of 2.2 percent per year, which is the value estimated by the World Bank for Ghana. The next three targets are the secondary school completion rate in the aggregate – meaning for all individuals of school age – and the secondary school completion rates in the top and bottom test score quartiles of the control group. The final non-experimental moment we target is the variance of the permanent component of log wages.[3] We target a value of 0.22 from the estimate of Lagakos and Waugh (2013), which is in line with other estimates found in the literature.

The remaining five moments come primarily from the experiment of Duflo, Dupas, and Kremer (2021), described in Section 2. To match these moments, we need to be able to replicate their experiment within our model. We describe how we do this in the following subsection. The last moment we target is the intergenerational correlation of schooling in Ghana, taken from Azomahou and Yitbarek (2021), and computed from regressions of children's educational attainment on parents' educational attainment. We target these regression coefficients by running these same regressions in our model.

4.3. Running the Experiment in the Model

We replicate the experiment in partial equilibrium. Since the experiment affected just 2,064 students, we find it implausible that the experiment had any significant general equilibrium effects. We also abstract away from the difference between day schools, which are the subject of the experiment, and boarding schools, which may be of higher quality, since day schools are more likely to be the focus of secondary schooling expansions in the future.

Importantly, we mimic the sample selection in the experiment, which consisted of picking "smart kids from poor families." To match the requirement that students in the sample have passed the BECE, we choose a test score cutoff so that only the top 42 percent of students in the model pass, consistent with the actual BECE passing

[3] Along the balanced growth path of the model, this object corresponds to $Cov\left(\log y_{i,t}, \log y_{i,t-1}\right)$, where the subscript i denotes agent and t denotes time. To see this, note that (5) implies $\log y_{i,t} = \log \zeta_{i,t} + \log\left(w_{i,t} h_i\right)$. Since $\log \zeta_{i,t}$, the transitory component of $\log y_{i,t}$, is i.i.d. and is uncorrelated with the permanent component, $\log\left(w_{i,t} h_i\right)$, we have $Cov\left(\log y_{i,t}, \log y_{i,t-1}\right) = Var\left(\log\left(w_{i,t} h_i\right)\right)$.

rate. Selection into the experiment also required that students had not registered for secondary school in the fall semester following their exam, which is harder to match literally within the model (particularly since one period in the model represents five years). Our strategy is to choose a parental income cutoff such that, among the students passing the BECE in the model, the eventual secondary school completion rate for those below the cutoff is 47.5 percent, just as in the control group of the experiment. We then choose the experimental sample in our model to be a subset of those with test scores *above* the test score cutoff and income *below* the income cutoff.

We treat the experiment as unanticipated and assume that model households know that the experiment ends after a single generation. Households selected into the control group solve their optimization problem as usual. Households selected into treatment experience an exogenous reduction in the goods cost of secondary school Ψ_S to 0 for the current period and then re-optimize. We construct simulated equivalents of the experimental moments by taking simple differences of average outcomes between the treated and control households in the model, which correspond to the intent-to-treat estimates in the experiment.

We target the negative treatment effects on fertility and positive treatment effects on human capital in the experiment, which we view as the most important – and precisely estimated – findings of the experiment. The treatment effect on fertility is large, and consistent with a 10.6 percent reduction in fertility after 12 years. The experiment found substantial positive impacts on test scores in reading and math of 0.16 standard deviations, which are consistent with the impacts of other successful interventions found in this literature (e.g. Duflo, Hanna, and Ryan, 2012; Mbiti, Muralidharan, Romero, Schipper, Manda, and Rajani, 2019).

Regarding their estimates of the labor market impact of the scholarship, Duflo et al. (2021) say "the total impact on earnings is 37 shilling (3% of the control group mean), a very imprecise estimate (95% CI [-10%,+15% of the control group mean], p-value 0.65). We cannot reject that returns are either zero or high compared to standard estimates of Mincerian returns." Despite heroic efforts to follow the RCT participants for around a decade, the estimated monetary returns to schooling are imprecise and cannot rule out anything ranging from substantial reductions in earnings to Micerian returns larger than those typically estimated in advanced countries like Denmark (see e.g. Schoellman, 2012).

For this reason, we instead choose to target the treatment effect on test scores as the experimental measure of the impact of receiving the scholarship on human capital. The estimate impact is 0.16 standard deviations and, in the model, we convert this increase in test scores to an increase in wages by assuming that a 0.16 standard deviation increase in test scores for the treatment group relative to the control group corresponds to a 0.16 standard deviation increase in wages for the treatment group, equivalent to wage gains of 7.6 percent in the quantitative model. This is higher than the point estimate for earnings of 3 percent from the experiment but is well within the confidence interval. An additional advantage of this approach is that targeting a direct measure of human capital gains such as test scores circumvents the concern that increases in earnings may (at least partially) reflect increased access to rent-earning public sector jobs and thus overstate the extent to which education improves labor productivity.

The treatment effects on school attendance are also informative about the extent of misallocation in education. We hence target the experiment's treatment effect on school completion, which was 27 percent. Additionally, we target the treatment effect on secondary school completion in the top quartile of test scores relative to the bottom one. This difference is small, at 4 percent, meaning that the overall treatment effect on secondary school completion was not particularly skewed toward those with high test scores relative to those with low scores.

4.4. Model Fit and Validation

Table 2 reports the targeted moments and their values in the estimated model. We also report the 95-percent confidence intervals for the moments that we resample in the bootstrap procedure. The fit is good for most moments, but a bit off for several of them, it must be said. On the plus side, the model does well in matching the treatment effects on human capital (6.7 percent versus 7.6 percent in the data) and fertility (-11.6 percent versus -10.6 percent in the data). The population growth rate and the variance of the permanent component of income are matched more or less exactly, and the model's treatment effect on secondary school completion is only 3 percentage points higher for the top quartile of the test score distribution than the bottom quartile, which is close to the 4 percent in the data.

The model is less successful in matching the average secondary schooling completion rates (30 percent versus 34 percent in the data), and the model's completion rates are a bit too high in the top test quartile and a bit too low in the bottom test quartile. Overall, though, the model captures the slight increase in completion rates by test score quartiles in both the control and treatment groups (see Appendix Figure A.1). The treatment effect on secondary school completion is too low in the

Table 2: Targeted Moments and Model Predictions

Moments	Data	Model
Aggregate Population Growth	2.2	2.2
Aggregate SHS Completion Rate	34	30
SHS Completion, Q4 of Test (Control Group)	53	65
SHS Completion, Q1 of Test (Control Group)	41	35
Var(Permanent Component of Income)	0.22	0.22
Treatment Effect on Human Capital	7.6 $(3.2,\ 12)$	6.7
Treatment Effect on Fertility	-10.6 $(-20.8,\ -0.4)$	-11.6
Treatment Effect on SHS Completion	27.0 $(22.7,\ 31.3)$	21.3
Treatment Effect on SHS Completion, Q4 - Q1	4 $(0,\ 8)$	3
Intergenerational Schooling Correlation	0.45 $(0.43,\ 0.47)$	0.32

Note: This table reports the moments targeted in the estimation and their values in the data and in the model. The range reported below each moment in the bottom half of the table (below the line) is its 95 percent confidence interval.

model, and the same is true of the intergenerational schooling correlation. In Appendix C, we perform some robustness exercises and show that missing on these two moments is not important for our main conclusions.

The estimated parameter values, and their bootstrapped confidence intervals, are presented in Table 3. While there is certainly some uncertainty in the estimated values, the confidence intervals for each parameter are fairly reasonable, suggesting that the model is precisely estimated in a statistical sense.

The estimated parameters seem reasonable from an economic sense as well. The estimated fertility parameter ν_J implies that each less-educated family has 1+1.07 = 2.07 offspring, implying an average of 4.1 children in a less-educated family with two adults. Similarly, the estimate of ν_S implies that each more-educated family has 1+0.19 = 1.19 offspring, which corresponds to around 2.4 children per family with two adults. These predictions are quite similar to (non-targeted) averages from the Demographic and Health Surveys for Ghana, which show average fertility of

Table 3: Parameter Estimates and Confidence Intervals

Parameter	Description	Estimate (Confidence Interval)
ν_J	Fertility of Primary School Graduates	1.07 $(1.03,\ 1.17)$
ν_S	Fertility of Primary School Graduates	0.19 $(0.17,\ 0.21)$
η_S	Efficiency of Secondary School	5.66 $(4.39,\ 6.14)$
Ψ_S	Goods Cost of Secondary School	1.52 $(1.48,\ 1.71)$
b	Intergenerational Altruism Factor	2.26 $(2.1,\ 2.45)$
σ_ε	Std. Deviation of Exam Score Noise	0.92 $(0.89,\ 1.04)$
θ	Gumbel Scale Parameter of Taste Shock	0.42 $(0.39,\ 0.46)$
χ	Savings Wedge	0.09 $(0.09,\ 0.10)$
ρ	Persistence of Ability Process	0.79 $(0.77,\ 0.92)$
σ_υ	Std. Deviation of Ability Shock	0.36 $(0.34,\ 0.39)$

Note: This table reports the estimated parameters. The confidence interval is the 2.5th and 97.5th percentiles of 100 bootstrapped parameter estimates.

4.1 children for women with JHS only, and 2.6 children for those with secondary education or more.

The estimated efficiency of schooling, η_S, is hard to interpret directly but implies (with all the other parameters) an annual return to education of 7.9 percent per year for this experimental sample. This is generally in line with other estimates of returns to education in developing countries, and if anything is on the high side. Schoellman (2012), for example, estimates returns of around 4 percent in Ghana and values generally under 5 percent for Sub-Saharan Africa (with large confidence intervals). The cost of schooling, Ψ_S, amounts to 25 percent of GDP per capita, which is close to the 21 percent reported by Duflo et al. (2021).

To better understand how plausible the estimate of b is, we compute the compensating variation of secondary schooling for all children in the model, at age 15 when their schooling decision is being made. We find that the average compensating variation is similar to the average cost of schooling (including opportunity cost), modestly lower for children who receive only a JHS education, and substantially larger for those whose parents choose a secondary education (see Appendix Figure A.2). These calculations imply that the children's valuation of schooling is mostly in line with that of their parents, suggesting that the value of b is reasonable.

The savings wedge, χ, has a value of 0.09, which implies that households save at around 11 percent per period, or 2 percent per year. This is a low return to savings but not as low as the negative returns posited by other similar incomplete-markets models estimated to data from developing countries (e.g. Lagakos, Mobarak, and Waugh, 2023; Donovan, 2021).[4]

The estimated value for the intergenerational persistence of ability, ρ, is 0.79, implying a strong correlation between parents' and children's ability. This is broadly consistent with the recent conclusions of Lee and Seshadri (2019) that parental traits, summarized by ability in our model, explain a substantial amount of the variation in children's income levels. The estimated standard deviation of the ability shock, σ_v, is 0.36. While not directly interpretable, this value (along with the other parameters) generates a Gini coefficient within the model of 0.31. This is somewhat lower than the Ghanaian value of 0.43 but within the range of 0.3 to 0.6 reported in the World Development Indicators for other Sub-Saharan African countries.[5]

5. Simulating the Effects of Free Secondary School

Using the estimated model, we simulate the effects of a national free secondary schooling policy. To do so we make two specific changes to our model. First, we

[4]The average household asset-to-income ratio in the model is 0.5. This is broadly in line with other estimates from low-income countries, such as Samphantharak and Townsend (2018), who find a ratio of around 0.6 in Thai villages. Unfortunately, we know of no reliable household asset data in Ghana to which we can make a direct comparison.

[5]In Appendix C we discuss how we view each parameter as being identified from the targeted moments. We do so by computing the elasticities of each model moment to each parameter, following Kaboski and Townsend (2011), and by computing the sensitivity of each parameter to each targeted moment, following Andrews, Gentzkow, and Shapiro (2017). In short, the population growth parameters, ν_J and ν_S, largely shape the model's aggregate population growth rate and treatment effects on fertility. The variance and persistence parameters of the ability process, σ_v and ρ, relate most closely to the permanent component of income variance and the intergenerational schooling correlation, though they affect other moments significantly as well. The efficiency of schooling, η_S, and the intergenerational altruism parameter, b, drive the benefits of schooling, and play a big role in the model's aggregate secondary attendance; η_S affects the treatment effect on human capital more directly. The cost of schooling, Ψ_S, reduces school attendance, raises the treatment effect on schooling and lowers the treatment effect on fertility. The savings wedge, χ, and the standard deviations of the test score noise and taste shocks all impact secondary attendance in the top versus bottom quartiles of the test score distribution.

eliminate the goods cost of schooling for parents, and force the government to pay for all the goods costs through tax revenues. Second, we eliminate the test score cut-off required for SHS admission. This is consistent with Ghana's actual free schooling policy. We assume that households do not anticipate the policy and that the economy is on the balanced growth path at the time of implementation.

Public funding of SHS in the model works as follows. We require that the government raise taxes in proportion to the existing tax rates. Before the policy, each household paid taxes according to the tax function T that is a function of parent's and child's income. The post-policy tax function takes the form $(1 + \tau_t)T$ where τ_t is the proportional increase in taxes each period. Taking this approach maintains the current structure of the labor tax schedule, and in particular, the feature that the poorest half of households pay no taxes (see Appendix Table A.3).

We choose τ_t so that per period tax revenue along the post-policy balanced growth path is equal to per period tax revenue along the pre-policy balanced growth path plus the additional education expenditure. In other words, we assume that the policy does not change per capita spending on public goods G.[6]

5.1. Quantitative Results

The general equilibrium effects of the policy are summarized in Table 4. We also report confidence intervals for each aggregate outcome using the bootstrapped parameter estimates summarized above. While this is a natural use of bootstrapped parameter estimates, it is not commonly done in G.E. counterfactual simulations in macroeconomics. The goal is simply to quantify the uncertainty in the model's predictions arising from sampling uncertainty in the targeted moments — the experimental moments in particular as these are estimated with large standard errors.

The number of secondary schooling graduates increases by about 12 percentage points, under free schooling, from 30 percent of the population to 42 percent. The increase is small relative to the changes in secondary school completion in the experiment, in large part because the experimental sample is highly selected relative to the general population. Fertility falls due to the schooling expansion, and the population growth rates fall by 0.2 percentage points. Adult earnings increase by about 1.2 percent from the policy, stemming largely from the higher wages for the 12 percent of the population now receiving secondary education. On the other hand, child

[6] We have experimented with alternative public finance arrangements but find that they make no substantive difference in our conclusions. For this reason we stick with the simpler assumption of period-by-period budget balance.

earnings decrease by 7.5 percent, representing the opportunity cost of the newly educated workers.

Table 4: G.E. Effects of Free Secondary Schooling

Statistic	Change Under Free Schooling
GDP per Capita (%)	-1.0 (-4.2, -0.7)
Secondary School Completion (p.p.)	11.8 (4.2, 13.1)
Population Growth Rate (p.p.)	-0.2 (-0.3, -0.1)
Adult Earnings (%)	1.2 (-4.4, 1.8)
Child Earnings (%)	-7.5 (-9.2, -1.5)
Taxes per Capita (%)	1.4 (1.1, 1.7)
Skilled Wage/Unskilled Wage (%)	-9.8 (-11.8, -6.5)
Gini Coefficient	-0.04 (-0.05, -0.04)

Note: This table reports the estimated aggregate effects of free secondary schooling. The changes in secondary school completion rates and population growth rates are expressed in percentage points. The changes in the Gini coefficient is measured in levels. The changes in all other statistics are expressed in percentage changes. The range reported below each estimated value is its bootstrapped 95 percent confidence interval for the change.

Rather than increasing, GDP per capita falls by about 1 percent in the long-run. Our confidence interval excludes decreases smaller than 0.7 percent in magnitude. Thus, from the perspective of sampling uncertainty in the targeted moments, we can reject any positive impact on GDP. The long-run cost of the policy is 1.4 percent of GDP, implying a total cost of the program of about 2.4 percent of GDP (the direct cost plus lost GDP). Relative wages of the skilled fall by about 10 percent, pointing to clear distributional impacts of free schooling policies, even for those who remain unskilled after the policy change. Our predictions here are similar at least qualitatively to those of Khanna (2023), who finds substantial declines in the relative wages of skilled workers after an education expansion in India. His wage effects are larger than ours quantitatively, though his study focuses on the short run where elasticities of substitution between low and high skilled workers are likely smaller.

Figure 3: Child SHS Attendance Probability in Estimated Model

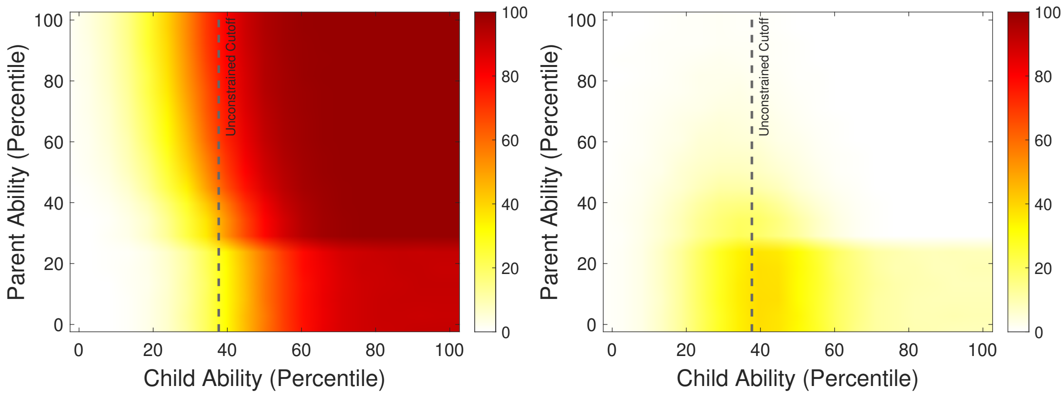

(a) Attendance Probability

(b) Increase in Attendance

The modest declines in GDP suggest that the estimated model does not feature high levels of misallocation in education. Recall that Figure 2 provided examples of economies with high and low misallocation by displaying the probability of SHS attendance as a function of child and parent ability. Figure 3, Panel (a), displays an identical plot using the fully estimated model. While it falls between the two extremes shown in Figure 2, it is clear that the estimated model corresponds more closely to the case with low misallocation; the highest ability children are fairly likely to attend SHS even if they are born to low ability parents. Attendance does have some dependence on parent ability; however, this effect is fairly minimal and the probability of attendance looks much closer to that of Panel (a) in Figure 2. Increases in attendance from free secondary schooling (Panel b) also more closely resemble those of the low-misallocation economy.

5.2. Opportunity Cost, Selection, and Differential Fertility

Given that the back-of-the-envelope calculation at the end of Section 2 suggested a sizable large GDP gain of around 6 percent based on the experimental estimates, it is natural to wonder what mechanisms lead the model to instead predict a modest reduction in GDP as a result of expanding secondary schooling. Here we perform counterfactual experiments using the model to highlight the role of three channels that each explain some portion of the difference.

We present the results of counterfactuals, which we describe below, in Figure 4. In each panel, the left-most brown bars reproduce the predictions for free secondary schooling coming from the baseline economy. The other four sets of bars, to the right of the baseline, represent the results of the four different counterfactual exercises. The panels represent the increases in GDP gains (top left), secondary completion rates (top right), adult earnings (bottom left), and child earnings (bottom right) coming from free secondary schooling.

First, we highlight the role of opportunity cost. To do this, we consider a counterfactual world in which the opportunity cost of secondary school is zero and simulate the effects of the free schooling policy in this world. We implement this in the model by simply eliminating the option for children to spend ages 15 through 20 working, so that parents' choice is between education and idleness (which neither the child nor the parent value) rather than education and income. We then recompute the pre-policy steady state of the model and simulate the effects of the free SHS in this new world.

The results of this exercise can be seen in the beige columns of Figure 4, marked 'no opportunity cost.' In this counterfactual, GDP per capita increases by 2.9 percent compared to the -1.0 percent drop in the baseline. This larger income gain is driven in part by substantially larger increases in secondary attendance of 23 percentage points, compared to 11.8 percentage points in the baseline. This larger increase in secondary completion leads to a larger increase in adult earnings (4.8 percent), and the lack of work during secondary years eliminates almost all of the decline in children's earnings. Overall, this counterfactual suggests that opportunity cost reduces the gains from free schooling by almost 4 percentage points (2.9 percent minus -1.0 percent).

The second counterfactual examines the role of selection stemming from the non-representativeness of the experimental sample. We simulate the effects of free schooling while simultaneously boosting the abilities of new secondary attendees so that they draw from the ability distribution of the experimental sample. More precisely, we start from the pre-expansion steady state and implement the free secondary school policy as in the baseline; however, every child who would not have gone to school according to the pre-expansion policy functions redraws their ability z randomly from the ability distribution of the experimental sample. After seeing the result of this draw, the parent decides whether or not to send the child to school. As a result, the set of new attendees looks identical to the experimental sample, effectively eliminating the role of selection.

The results are displayed as the orange bars marked 'No selection on ability'

Figure 4: Effects of Free SHS under Counterfactual Scenarios

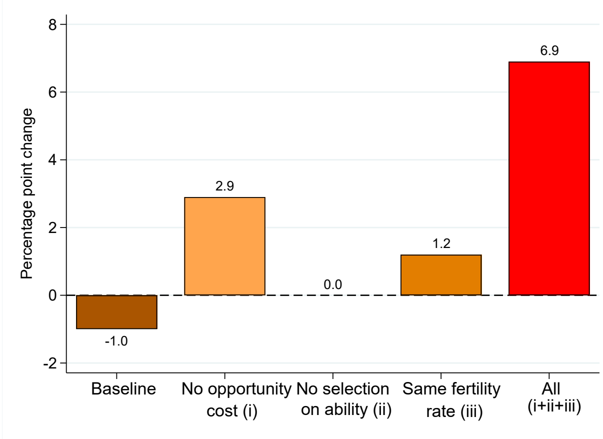

(a) GDP per Capita

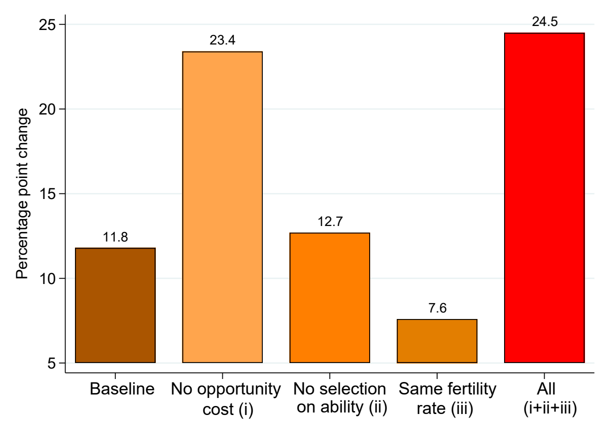

(b) SHS Completion Rate

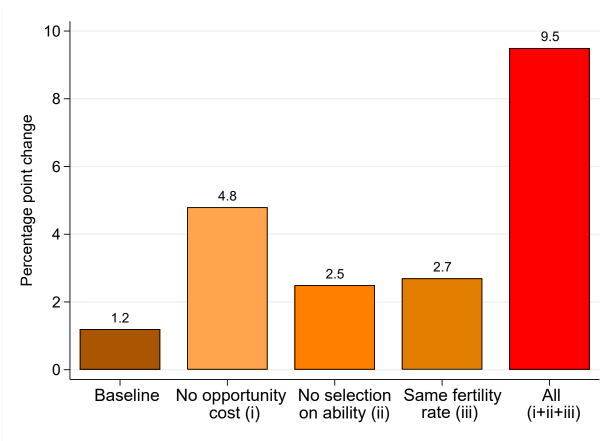

(c) Adult Earnings

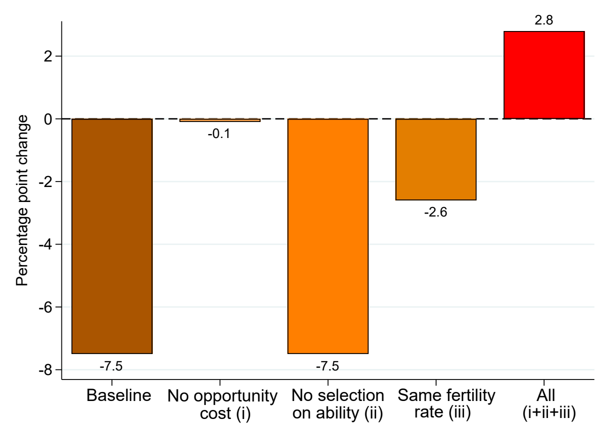

(d) Child Earnings

in Figure 4. Rather than the 1 percent decrease predicted by the baseline model, GDP per capita is essentially unchanged in this counterfactual. The reason is that earnings of adults now increase by 2.5 percent – around twice the increase observed in the baseline. In other words, the fact that the experimental sample exhibits higher average ability than the general population knocks off around a percentage point from the GDP gains from expanding secondary access.

We investigate the role of fertility changes in the third counterfactual. Like with the selection counterfactual, we begin from the pre-expansion steady state of the model, and then implement the free schooling policy. We then change the fertility level of new secondary attendees to be $1 + \nu_J$, the fertility level of those without

Figure 5: Changes in Ability Distribution Induced by Free Schooling

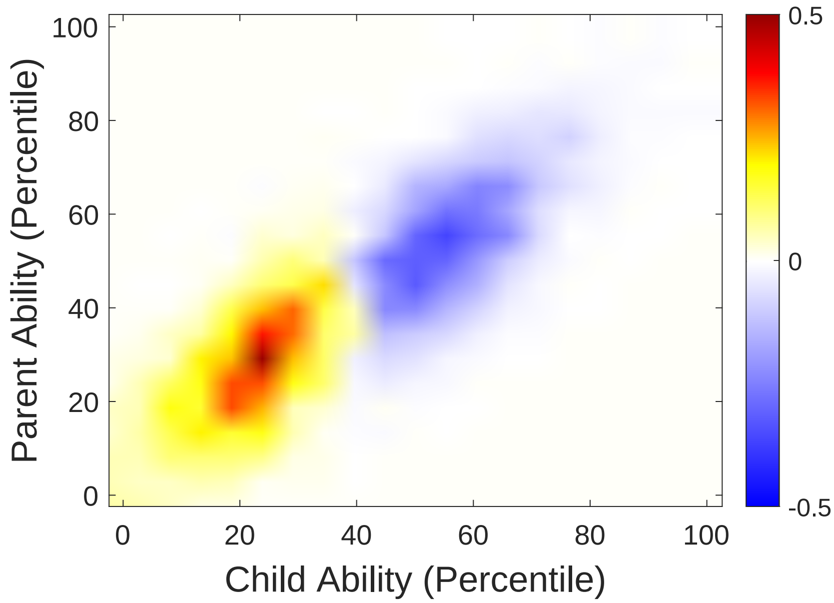

secondary education. This counterfactual therefore allows free schooling to operate as before but now without affecting fertility for new attendees.

The dark orange bars labeled ’Same fertility rate’ report the results of this exercise. The GDP per capita gains are now modestly positive at 1.2 percent. Like in the case of selection, this is largely driven by an increase in adult earnings, which expand by 2.7 percent. Secondary attendance increases by somewhat less than the baseline, and child earnings decrease by less. To see why adult earnings increase once the fertility channel is shut off, Figure 5 plots the percentage point changes in the (post-policy) distribution of ability in this case relative to the baseline. The blue portions of the figure are where population declines (by around 0.5 percent of the population) and the red portions indicate places where population increases (also by around 0.5 percent). The magnitudes are not dramatic, but the overall patterns are. There is a significant shift from higher ability parents and children to lower-ability ones. This figure highlights how unfavorable long-run movements in the ability distribution serve to reduce the gains from free schooling by about 2.2 percentage points.

The red (rightmost) bars of Figure 4 display what happens when we combine all three counterfactuals, shutting down the opportunity cost, selection and differential fertility channels. In this case the GDP gains from free schooling are large, at 6.9 percent, and close to 8 percentage points higher than in the baseline. Now secondary schooling expands by a robust 24.5 percentage points and adult earnings grow by close to 10 percent, while child earnings also grow modestly. Unlike in the main analysis, taxes per capita stay roughly the same in this counterfactual as free schooling largely pays for itself. Note that the overall GDP gain in this counterfactual is close to the back-of-the-envelope gains we reported in Section 2 taking into account just the returns to schooling proxied by test scores and the increase in years of schooling.

A key takeaway from these counterfactuals is that the model has no problem delivering a sizable gain in GDP from free schooling once we shut down its opportunity cost, selection and differential fertility channels. Thus, we conclude that these forces are the key ones that combine to eliminate nearly all the gains from higher human capital from free schooling implied by the significant test score and secondary completion increases found by Duflo, Dupas, and Kremer (2021).[7]

5.3. Welfare and Distributional Impacts

We turn next to the welfare and distributional impacts of the free secondary school. Welfare is inherently hard to discuss in this setting due to the presence of endogenous fertility. Any policy that changes schooling attendance decisions also changes fertility, and the set of agents that exists in the post-policy balanced growth path is not the same as the set that would have existed had the policy never been implemented. In our case, free secondary school leads to higher attendance and lower fertility. Thus some individuals who would have existed had the policy not been implemented are instead never born.

Rather than taking a stance on how to aggregate welfare across agents who may or may not exist after a policy change (as in Golosov, Jones, and Tertilt, 2007; De la Croix and Doepke, 2021), we simply report consumption-equivalent welfare separately for different groups of agents. In particular, we focus on parents with newborns at the time of policy implementation and their eventual children. Because these parents have already given birth, their fertility decisions are determined, and all agents in these two groups exist both with and without the policy, making traditional welfare comparisons possible. Additionally, to get a sense for how welfare and fertility changes interact, we examine welfare for two groups of the parents'

[7] The results of this section highlight that our model's negative predictions about GDP per capita are not driven by the fact that the experiment's treatment effect on labor earnings is wide and includes zero. We do not target this moment in our estimation.

Table 5: Welfare Change of Select Groups under Free Schooling

	Overall	Bottom 25 Pct	Top 25 Pct
Parents	3.9	7.3	-4.2
Children	12.4	23.1	-5.3
Grandchildren (always born)	9.0	-	-
Grandchildren (born w/o policy)	16.9	-	-

Note: This table reports the average change in consumption-equivalent welfare for select groups of individuals under the free schooling policy.

grandchildren: those who would always exist regardless of schooling policy and those who would exist only if schooling were not expanded.

For each of these groups, we compute lifetime utility from consumption, ignoring utility gained from altruism towards children (or grandchildren or great-grandchildren, etc.). We compute this value under both the free secondary schooling policy and the case of no policy change, and our measure of consumption-equivalent welfare reports the percentage increase in the consumption of all individuals (within the relevant group) required to raise the average utility level under no policy change to that achieved by the policy.

We are also interested in the redistributive component of the policy; that is, how much of the welfare gains accrue to poor households relative to rich ones. As in Fernández and Rogerson (1995), rich households in our model are more likely to go to school. Thus a free SHS policy risks being regressive. Unlike Fernández and Rogerson (1995), who model schooling as funded through proportional taxation, this effect is mitigated by the fact that our tax schedule is strongly progressive. Thus the redistributive nature of the policy is a quantitative question. In order to answer this question, we also report the change in welfare for parents and children in the top and bottom 25 percent of the (pre-policy) income distribution.

The first two rows of Table 5 report our welfare measure for parents (those with newborn children at the time of policy implementation) and their children. Welfare for parents increases on average by 3.9 percent of consumption. This is somewhat surprising as, due to higher school attendance and the loss of children's wages, the policy actually reduces income in this group. Instead, the welfare gains come entirely from redistribution. The poorest 25 percent of parents gain 7.3 percent of consumption while the richest lose 4.2 percent. Despite concerns that wealthy parents are more likely to send their children to secondary school and thus are more likely to be beneficiaries of the policy, the tax schedule is sufficiently progressive that the policy acts redistributively overall.

The gains for children are large and positive at 12.4 percent. As was the case for parents, the policy is highly redistributive. The children who would have ended up in the top 75 percent of the income distribution absent the policy lose 5.3 percent, roughly the same amount as the parents, due to higher taxes. The (pre-policy) poorest 25 percent of children, however, make substantial gains of 23 percent. While a small portion of these gains occur broadly within this group due to higher unskilled wages, the majority are accrued by the relatively small number of misallocated children who make substantial income gains.

The welfare results discussed so far present a small puzzle. Despite a decline in long-run output per person of 1 percent, the policy seems to lead to large gains in welfare, particularly for children who gain more than can be plausibly attributed to redistribution. This tension is resolved by examining the welfare impacts in the longer run. The third and fourth rows of Table 5 list welfare gains for the grandchildren of the parents described in the first row, separated into those who will always be born regardless of policy and those who would only be born in the pre-policy steady state.[8] While the welfare gains for the grandchildren who exist regardless of policy are positive, they are smaller than those of the children (9 percent vs 12 percent). The central reason is that the grandchildren who do not get born under the free schooling policy experience disproportionately large benefits (almost 17 percent). Essentially, the decline in fertility serves to shift the composition of the population towards those who experience smaller gains. This does not stop after the grandchildren are born and continues to compound with each generation, eventually leading to the long-run decline in GDP per capita described in Table 4.

5.4. Robustness and Discussion

Overall, the results of this section suggest that the free secondary schooling policy does not lead to a long-run increase in overall income per capita and instead serves mostly as redistribution. In order to focus on the channels of opportunity cost, selection, and fertility, the model used to draw these conclusions simplifies away some potentially important features. Here, we discuss some of these features and perform

[8] Conceptually, the welfare for grandchildren who would not be born if the policy were implemented is computed by asking the question "if the policy were implemented and your dynasty alone was exogenously assigned to experience no fertility change (so that you, but none of the others in your category, exist), what would your change in welfare be?"

some basic quantitative exercises to assess the robustness of these conclusions.

Familial Bargaining and Transfers: One restriction placed upon the model is an abstraction away from household bargaining in favor of the assumption that parents dictate the decisions of children and that there are no transfers (in either direction) between parents and children. While this provides tractability, it raises the possibility that parents and children disagree on the value of schooling and lack a market in which to interact and resolve this disagreement. Further, the lack of non-schooling transfers means that secondary school is the only mechanism through which parents can make transfers towards children. These features raise the possibility that the conclusion that the gains from free schooling are small arises from parents over-valuing secondary school and "over-schooling" by enrolling children whose private valuation of school would not have them attend.

To investigate whether or not this is a concern, we perform a compensating variation calculation to assess children's valuation of schooling. In particular, we compute the level of assets at which a child becomes indifferent between being "born" (i.e. forming their own household) possessing a JHS education and the stock of assets or being born possessing an SHS education and no assets. In essence, this calculation computes the child's value of secondary education.

Appendix Figure A.2 displays the average compensating variation (as a percent of GDP per capita) for SHS and JHS graduates in the pre-policy steady-state of the model. Reassuringly, the average valuation of secondary schooling among those who receive it is much larger than the total (i.e. opportunity-cost-inclusive) cost of the schooling (217 percent of GDP per capita vs 83 percent). Similarly, the average valuation for those who do not attend secondary school is smaller than, but close in magnitude to, the cost (78 percent of GDP per capita vs 83). At least on average, children who value secondary schooling above its cost are attending while those who value it under the cost are not, assuaging concerns that the lack of familial bargaining or transfers may be driving the headline results.

Human Capital Externalities: Externalities through which increases in the stock of skills improve overall productivity, leading the aggregate gains in output due to an increase in education to be larger than the sum of the private gains (e.g. Acemoglu and Angrist, 2000), are often referenced as reasons for education expansion and are completely absent from the baseline model. The absence of such human capital externalities would lead the model to underpredict the increase in GDP per capita arising from an increase in educational attainment due to the free secondary school policy.

Fortunately, extending the model to include such an effect is fairly simple; adding a term $H_{\text{ext}}^{\tilde{\beta}}$ to the production function yields

$$Y = AH_{\text{ext}}^{\tilde{\beta}}K^{\alpha}\left[(N_J)^{\lambda} + (N_S)^{\lambda}\right]^{\frac{1-\alpha}{\lambda}} \tag{11}$$

which can then be combined with the equilibrium condition $H_{\text{ext}} = N_S/P$ to implement the human capital externality, ensuring the neither firms nor individuals internalize the social benefit of increasing the average skill level in the population.

With this production function, the ratio of the marginal social benefit of moving one individual from primary to secondary education (i.e. decreasing N_J by unity and increasing N_S by the individual's z) and the marginal private benefit does not depend on z and is given by $1 + \frac{\tilde{\beta}}{1-\alpha}\frac{N_J^{\lambda}+N_S^{\lambda}}{N_S^{\lambda}}$. This expression yields a correspondence between $\tilde{\beta}$ and the spillovers generated by the externality — if an individual's return to a year of education is 10 percent, but increasing the economy-wide level of education by one year increases output by 12 percent (suggesting a spillover of 20 percent), the corresponding value of $\tilde{\beta}$ is given by $(1-\alpha)(1.2-1)\frac{N_S^{\lambda}}{N_J^{\lambda}+N_S^{\lambda}}$.

We choose $\tilde{\beta}$ to yield a generous spillover of 50 percent (i.e. the social gains of education are 1.5 times larger than the private gains), corresponding to a value of about 0.16. This is substantially larger than many estimates in the literature, which often find a small role for such externalities (Acemoglu and Angrist, 2000; Ciccone and Peri, 2006), particularly in the case of secondary schooling (Iranzo and Peri, 2009), and even exceeds the larger estimates (e.g. Gennaioli, La Porta, Lopez-de Silanes, and Shleifer, 2013, who find a spillover of about 40 percent).[9]

Even after incorporating this large externality, implementing the secondary school expansion policy continues to yield results similar to those in the baseline model. Rather than decreasing by 1 percent, GDP now increases by 0.2 percent. Although it is now positive, the impact of the policy on GDP remains anemic and substantially smaller than the back-of-the-envelope calculation in Section 2 would suggest. The aggregate return human capital is larger than in the baseline model but, as discussed above, the estimated model does not feature particularly high levels of misallocation and, consequently, the schooling expansion policy only modestly increases aggregate human capital, leading to small results even when an externality is included.

[9]In particular, Gennaioli et al. (2013) estimate a Mincer return to schooling of 20 percent and a regional TFP increase from an additional year of education of 7.5 percent, suggest a spillover of 37.5 (7.5/20) percent

Rationed Public Sector Jobs: In many developing countries, the ability to obtain public sector jobs, which command a large wage premium relative to the private sector, is a large part of the reason that many choose to complete secondary education (or beyond). While these jobs lead to a substantial increase in earnings for the individuals who obtain them, the common view among economists is that this premium reflects government rents rather than a higher productivity of the public sector. Because these rents are not subject to market pressure and effectively act as a wage floor, these jobs end up being rationed.

It is unclear how much of an issue this is for our results. To the extent that higher earnings for secondary school graduates reflect these rents, rather than higher productivity, the model will overestimate the impact of free schooling. However, the model is estimated to match the increase in human capital, measured via test scores, in the treatment group rather than the increase in wages which is less subject to this concern. This increase in human capital increases output, even if government rents impact how that output is distributed across individuals.

Still, it is useful to extend the model to examine the potential magnitude of this issue. Duflo et al. (2021) report that 7.7 percent of their control group (secondary graduation rate of 42.3 percent) end up employed in a public sector job. Assuming that all public sector jobs are allocated to secondary graduates,[10] this implies roughly an 18 percent chance of obtaining public employment conditional on completing secondary school. Finan, Olken, and Pande (2017) estimate a public sector earnings premium (in Ghana) of 76 percent. Attributing all of this earnings premium to rents (rather than higher productivity) and combining these two numbers ($(1-0.177)+1.76(.177)-1=0.13$) suggests that, at most, public sector rents account for 13 percent of the average return to secondary schooling.

To account for this in the model, we augment the human capital function h so that the human capital for someone with SHS education becomes $z\eta_S\big(1-0.13(1-\frac{1}{\hat{\text{SHS}}})\big)$ where $S\hat{H}S$ is the percent deviation of the aggregate SHS share from the pre-policy steady-state. Essentially, if the number of secondary school graduates doubles ($S\hat{H}S=2$), the return to secondary schooling falls by 6.5 percent, reflecting that the fact that the probability of any particular student obtaining a public sector job (which accounted for 13 percent of the pre-policy return) has been cut in half.

In this extension of the model, the impact of the free secondary schooling ex-

[10]Violations of this assumption would imply that our estimate for the contribution of the public sector earnings premium to the total SHS earnings premium is an upper bound.

pansion policy is almost identical to that of the baseline model. Unsurprisingly, the decline in GDP is slightly larger than the baseline model at -1.8 percent, a result of the fact that this extension to the model unambiguously shrinks the potential output gains from secondary school. The fact that the decline is only slightly larger than the baseline model reflects the fact that the increase in secondary attendance is modest, only about 35 percent (11 percentage point) and, consequently, the decline in the return to schooling induced by the model extension is small (about 5 percent).

6. Aggregate Effects of Alternative Policies

If making secondary school free to parents does not raise living standards, are there any alternative policy levers that governments in low income countries can pull to bolster their education systems and raise their average income levels? In this section we address this question by using our estimated model to simulate the aggregate effects of some alternative education policies.

Table 6: Aggregate Effects of Alternative Policies

	GDP Gain (%)	SHS Increase (p.p.)
Free Secondary School (Main Analysis)	-1.0	11.8
Free Secondary School + Relaxed Test	-0.3	6.3
Free Secondary School + Keep Test As-Is	0.1	3.0
Only Remove Test	-0.4	8.4
Tax + Transfer to Bottom 25%	0.1	-1.6
Raise Schooling Quality	2.7	13.8

Note: This table reports the gains in GDP and the increase in the SHS graduation rate under free schooling and several alternative policies (described in the text).

Table 6 summarizes the results of various alternative policy counterfactuals. The first row reproduces the key aggregate statistics from the free schooling policy counterfactual from the previous section (i.e. the gains in GDP and the increase in secondary schooling completion). The second reports the same outcomes in an alternative simulation where free secondary schooling is offered but, unlike the main analysis, the requirement that secondary students pass an admissions test is maintained. This policy mitigates GDP losses by maintaining higher standards for who can enter secondary schooling, reducing the problems of negative selection found in the main analysis. GDP per capita now only falls by 0.3 percent and the change in secondary attendance is smaller. The third row performs a similar exercise but does not change the stringency of the test at all, leading to basically no change in GDP per capita. On the other hand, the fourth row reports the effects of eliminating the test requirement but maintaining a positive cost of schooling. As in the main analysis, this leads to an increase in attendance but a decline in GDP.

The fifth row simulates a pure tax-and-transfer scheme to isolate the redistributive component of the free secondary schooling policy. We increase the tax rates by the same magnitudes as in the baseline policy but now simply distribute the extra revenues as a lump-sum transfer to the bottom 25 percent of earners. SHS decreases slightly as the increase in taxes distorts parents' educational decisions for their children; however, GDP is left essentially unchanged (as the students pushed out of SHS have marginal opportunity-cost-inclusive return to schooling).

The last row of Table 6 summarizes the effects of improving schooling quality in such a way that average test scores rise by 0.1 standard deviations and eliminating the admissions test. This effect is conservative relative to the average effect size estimated in a number of different randomized interventions aimed at improving schooling quality in the developing world, many of which find effects of around 0.2 standard deviations or higher. One such intervention is to offer financial incentives to teachers based on the test scores of their students. Muralidharan and Sundararaman (2011) and Duflo et al. (2012) found that this raised test scores in India for example, while Mbiti et al. (2019) found effects of a similar size for teacher incentives plus block grants for schools in Tanzania. Another successful schooling quality intervention is to increase the number of teachers in the classroom, as in the studies of Banerjee et al. (2007) and Muralidharan and Sundararaman (2013) in India. For our simulated intervention, we use the policy cost from Mbiti et al. (2019) who report that the cost of increasing test scores by 0.1 standard deviations per student in Tanzania was US $5.78.

Our model implies larger effects on GDP and welfare of improving schooling quality than providing free schooling. GDP increases by 2.7 percent under such an intervention. Even though this policy has no provisions aimed at expanding secondary enrollment directly, improved schooling quality raises school enrollments by 13.8 percentage points, slightly more than the free schooling policy. The implication is that many students were not attending secondary schooling to begin with because they felt the returns were not high enough to justify the costs (including

Table 7: Welfare Change of Select Groups under Improved Schooling Quality

	Overall	Bottom 25 Pct	Top 25 Pct
Parents	3.0	4.8	-0.8
Children	18.6	24.4	6.9
Grandchildren (always born)	14.5	-	-
Grandchildren (born w/o policy)	26.4	-	-

Note: This table reports the average change in consumption-equivalent welfare for select groups of individuals when schooling quality is improved.

opportunity cost).

These gains to GDP suggest that the channels emphasized in Section 5 are not so large as to prevent any increases in output from education interventions. While opportunity cost is still present, schooling quality improvements amplify the benefits of attendance directly and increase incomes even for inframarginal students, leading to larger output gains per "unit" of opportunity cost paid. Quality improvements also mitigate the impact of the selection channel; although new attendees may still have lower average ability than the experimental sample or current attendees, the fact that higher schooling quality affects both new and current attendees means that human capital increases even for the highest-ability students .

Table 7 lists the changes in consumption-equivalent welfare under improved schooling quality for the same groups of individuals as the discussion of welfare under the baseline policy in Section 5.3. Unsurprisingly, the gains in welfare for parents are smaller under the policy of improved quality (3.0 percent vs 3.9 percent) as, unlike the free secondary schooling policy, it does not represent a direct transfer of resources to parents. Instead, the gains in welfare are largely due to redistribution that occurs due to the change in relative wages.

Children and grandchildren benefit substantially more under improved schooling quality (18.6 percent vs 12.4 percent for children). This difference is almost entirely attributable to improvements in the welfare gains for the richest 25 percent of children who go from losing 5.3 percent under the baseline policy to gaining 6.9 percent when school quality is improved. The reason for this substantial increase is twofold. First, unlike the baseline policy of free schooling which improves the income of only marginal attenders, improvements in school quality also improve outcomes for inframarginal attenders (who make up almost all the top 25 percent).

Second, the policy is highly cost effective and pays for itself in the long run, resulting in a reduction in taxes — a benefit that mostly accrues to the richest. Despite this substantial improvement for the richest 25 percent, the poorest 25 percent still do better than under the baseline policy (24.4 percent vs 23.1 percent), suggesting that these improvements do not come at the cost of welfare for the very poorest.

7. Conclusions

One of the main reasons income per capita is so low in the developing world is that human capital levels are so low (Hall and Jones, 1999; Bils and Klenow, 2000; Erosa et al., 2010; Manuelli and Seshadri, 2014; Hendricks and Schoellman, 2018). Since attendance rates in primary schools are high in most low income countries, attention has turned to increasing secondary education. Making secondary schooling free for students, and funding the costs through higher taxes, is a natural policy option to consider. Not surprisingly, many developing countries are currently considering or implementing free secondary schooling policies of some kind.

In this paper we analyze the macroeconomic effects of free secondary schooling policies in the developing world, looking through the lens of an OLG model of human capital accumulation with credit constraints. We focus on the case of Ghana, for which we can draw on recent experimental evidence on the outcomes of students randomly assigned to receive free secondary schooling, leading to higher secondary school completion rates and higher average test scores (Duflo et al., 2021). Ghana is also a country that has recently adopted free secondary schooling, and the policy is viewed as a success there and in other developing countries (Center for Global Development, 2022).

Our conclusions are less optimistic. When we simulate the general equilibrium effects of free secondary school in our model, we find that it raises secondary enrollment substantially but has a negative impact on GDP per capita. Three forces in the model serve to offset the human capital gains from the policy. The first is lost earnings of secondary aged individuals, who forgo work in order to attend school. The second is the worse selection of students by ability level in school, which lowers the average effect of secondary education. The third is relative declines in fertility for the high ability, whose share in the overall population declines.

Of course, there are benefits of education that are beyond the scope of our model. For example, education expansions have been shown to reduce crime (e.g. Lochner and Moretti, 2004), create more informed voters, and improve parental ability (e.g. Daruich, 2020). We abstract from these channels largely due to a lack of evidence for our setting (though Duflo et al. (2021) found no evidence that increased school attendance altered voting behavior) and because we view them as secondary relative to the direct impact of human capital on productivity. Moreover, we conjecture that given our modest estimated effects of free schooling policies on schooling completion, adding these effects would be unlikely to have much additional impact.

We conclude that free secondary education policies are mostly redistributive in nature, rather than a path to economic growth, at least at current low levels of schooling quality. Improving schooling quality would lead to substantially higher GDP per capita and welfare across the income distribution for less fiscal cost. Improving schooling quality would also expand schooling enrollments by around the same amount as free schooling policies. Our analysis suggests that the best way for poor countries to raise human capital levels is to focus on improving the quality of their current schooling systems rather than giving away a mediocre education to more young people.

References

ABBOTT, B., G. GALLIPOLI, C. MEGHIR, AND G. L. VIOLANTE (2019): "Education Policy and Intergenerational Transfers in Equilibrium," *Journal of Political Economy*, 127, 2569–2624.

ACEMOGLU, D. AND J. D. ANGRIST (2000): "How Large are Human Capital Externalities? Evidence from Compulsory Schooling Laws," *NBER Macroeconomics Annual 2000*, 9–59.

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659–84.

AKCIGIT, U., H. ALP, AND M. PETERS (2021): "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries," *American Economic Review*, 111, 231–75.

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 132, 1553–1592.

ATTANASIO, O. P., C. MEGHIR, AND A. SANTIAGO (2012): "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 79, 37–66.

AZOMAHOU, T. T. AND E. YITBAREK (2021): "Intergenerational Mobility in Education: Is Africa Different?" *Contemporary Economic Policy*, 39, 503–523.

BANERJEE, A. V., S. COLE, E. DUFLO, AND L. LINDEN (2007): "Remedying Education: Evidence from Two Randomized Experiments in India," *Quarterly Journal of Economics*, 122, 1235–1264.

BASSI, V., R. MUOIO, T. PORZIO, R. SEN, AND E. TUGUME (2022): "Achieving Scale Collectively," *Econometrica*, 90, 2937–2978.

BECKER, G. S. AND N. TOMES (1979): "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility," *Journal of Political Economy*, 87, 1153–89.

BÉNABOU, R. (2002): "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70, 481–517.

Bewley, T. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252–292.

Bils, M., B. Kaymak, and K.-J. Wu (2022): "Labor Substitutability among Schooling Groups," NBER Working Paper 29895.

Bils, M. and P. J. Klenow (2000): "Does Schooling Cause Growth?" *American Economic Review*, 90, 1160–83.

Brooks, W. and K. Donovan (2020): "Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua," *Econometrica*, 88, 1965–1997.

Brudevold-Newman, A. (2021): "Expanding Access to Secondary Education: Evidence from a Fee Reduction and Capacity Expansion Policy in Kenya," *Economics of Education Review*, 83.

Buera, F. J., J. P. Kaboski, R. Rogerson, and J. I. Vizcaino (2022): "Skill-Biased Structural Change," *Review of Economic Studies*, 89, 592–625.

Buera, F. J., J. P. Kaboski, and Y. Shin (2021): "The Macroeconomics of Microfinance," *Review of Economic Studies*, 88, 126–161.

Buera, F. J., J. P. Kaboski, and R. M. Townsend (2023): "From Micro to Macro Development," *Journal of Economic Literature*, 61, 471–503.

Caunedo, J. and N. Kala (2022): "Mechanizing Agriculture," Unpublished Working Paper, MIT.

Center for Global Development (2022): *Schooling for All: Feasible Strategies to Achieve Universal Education*, Center for Global Development.

Chatterjee, S. and T. Vogl (2018): "Escaping Malthus: Economic Growth and Fertility Change in the Developing World," *American Economic Review*, 108, 1440–1467.

Ciccone, A. and G. Peri (2006): "Identifying Human-Capital Externalities: Theory with Applications," *Review of Economic Studies*, 73, 381–412.

Cunha, F. and J. Heckman (2007): "The Technology of Skill Formation," *American Economic Review Papers and Proceedings*, 97, 31–47.

DARUICH, D. (2020): "The Macroeconomic Consequences of Early Childhood Development Policies," Unpublished Working Paper, University of Southern California.

DE LA CROIX, D. AND M. DOEPKE (2021): "A soul's view of the optimal population problem," *Mathematical Social Sciences*, 112, 98–108.

DOEPKE, M. AND F. ZILIBOTTI (2017): "Parenting with Style: Altruism and Paternalism in Intergenerational Preference Transmission," *Econometrica*, 85, 1331–1371.

DONOVAN, K. (2021): "The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity," *Review of Economics Studies*, 88, 2275–2307.

DUFLO, E., P. DUPAS, AND M. KREMER (2021): "The Impact of Free Secondary Education: Experimental Evidence from Ghana," Unpublished Working Paper, Stanford University.

DUFLO, E., R. HANNA, AND S. P. RYAN (2012): "Incentives Work: Getting Teachers to Come to School," *American Economic Review*, 102, 1241–1278.

EROSA, A., T. KORESHKOVA, AND D. RESTUCCIA (2010): "How Important Is Human Capital? A Quantitative Theory Assessment of World Income Inequality," *Review of Economic Studies*, 77, 1421–1449.

FERNÁNDEZ, R. AND R. ROGERSON (1995): "On the Political Economy of Education Subsidies," *Review of Economic Studies*, 62, 249–62.

FILMER, D. AND N. SCHADY (2014): "The Medium-Term Effects of Scholarships in a Low-Income Country," *Journal of Human Resources*, 49, 663–694.

FINAN, F., B. A. OLKEN, AND R. PANDE (2017): "The personnel economics of the developing state," *Handbook of economic field experiments*, 2, 467–514.

GALOR, O. AND J. ZEIRA (1993): "Income Distribution and Macroeconomics," *Review of Economic Studies*, 60, 35–52.

GENNAIOLI, N., R. LA PORTA, F. LOPEZ-DE SILANES, AND A. SHLEIFER (2013): "Human capital and regional development," *The Quarterly journal of economics*, 128, 105–164.

GOLOSOV, M., L. E. JONES, AND M. TERTILT (2007): "Efficiency with endogenous population growth," *Econometrica*, 75, 1039–1071.

HALL, R. E. AND C. I. JONES (1999): "Why Do Some Countries Produce So Much More Output per Worker than Others?" *Quarterly Journal of Economics*, 114, 83–116.

HANUSHEK, E. A. AND L. WOESSMANN (2007): *Education Quality and Economic Growth*, World Bank.

——— (2012): "Do Better Schools Lead to More Growth? Cognitive Skills, Economic Outcomes, and Causation," *Journal of Economic Growth*, 17, 267–321.

HENDRICKS, L. AND T. SCHOELLMAN (2014): "Student Abilities During the Expansion of U.S. Education," *Journal of Monetary Economics*, 63, 19–36.

——— (2018): "Human Capital and Development Accounting: New Evidence from Wage Gains at Migration," *Quarterly Journal of Economics*, 133, 665–700.

HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 87, 1439–1474.

HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

IRANZO, S. AND G. PERI (2009): "Schooling externalities, technology, and productivity: Theory and evidence from US states," *The Review of Economics and Statistics*, 91, 420–431.

JENSEN, A. (2022): "Employment Structure and the Rise of the Modern Tax System," *American Economic Review*, 112, 213–32.

KABOSKI, J. P. AND R. M. TOWNSEND (2011): "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative," *Econometrica*, 79, 1357–1406.

KHANNA, G. (2023): "Large-Scale Education Reform in General Equilibrium: Regression Discontinuity Evidence from India," *Journal of Political Economy*, 131.

LAGAKOS, D., A. M. MOBARAK, AND M. E. WAUGH (2023): "The Welfare Effects of Encouraging Rural-Urban Migration," *Econometrica*, 91, 803–837.

Lagakos, D. and M. E. Waugh (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103, 948–980.

Laitner, J. (1997): "Intergenerational and Interhousehold Economic Links," *Handbook of Population and Family Economics*, 1, 189–238.

Lee, S. Y. and A. Seshadri (2019): "On the Intergenerational Transmission of Economic Status," *Journal of Political Economy*, 127, 855?921.

Lochner, L. and A. Monge-Naranjo (2011): "The Nature of Credit Constraints and Human Capital," *American Economic Review*, 101, 2487–2529.

Lochner, L. and E. Moretti (2004): "The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports," *American Economic Review*, 94, 155–189.

Loury, G. C. (1981): "Intergenerational Transfers and the Distribution of Earnings," *Econometrica*, 843–867.

Manuelli, R. E. and A. Seshadri (2014): "Human Capital and the Wealth of Nations," *American Economic Review*, 104, 2736–62.

Mbiti, I., K. Muralidharan, M. Romero, Y. Schipper, C. Manda, and R. Rajani (2019): "Inputs, Incentives, and Complementarities in Education: Experimental Evidence from Tanzania," *Quarterly Journal of Economics*, 134, 1627–1673.

Midrigan, V. and D. Xu (2014): "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, 104, 422–458.

Moll, B. (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review*, 104, 3186–3221.

Muralidharan, K. and V. Sundararaman (2011): "Teacher Performance Pay: Experimental Evidence from India," *Journal of Political Economy*, 119, 39–77.

——— (2013): "Contract Teachers: Experimental Evidence from India," NBER Working Paper 19440.

Porzio, T., F. Rossi, and G. Santangelo (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112, 2774–2814.

Restuccia, D. and C. Urrutia (2004): "Intergenerational Persistence of Earnings: The Role of Early and College Education," *American Economic Review*, 94, 1354–1378.

Samphantharak, K. and R. M. Townsend (2018): "Risk and Return in Village Economies," *American Economic Journal: Microeconomics*, 10, 1–40.

Schoellman, T. (2012): "Education Quality and Development Accounting," *Review of Economic Studies*, 79, 388–417.

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

Todd, P. E. and K. I. Wolpin (2006): "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96, 1384–1417.

——— (2023): "The Best of Both Worlds: Combining Randomized Controlled Trials with Structural Modeling," *Journal of Economic Literature*, 61, 41–85.

Appendix (for Online Publication)

A. Appendix Figures and Tables

Figure A.1: SHS Completion by Quartile of Test Score: Data vs Model

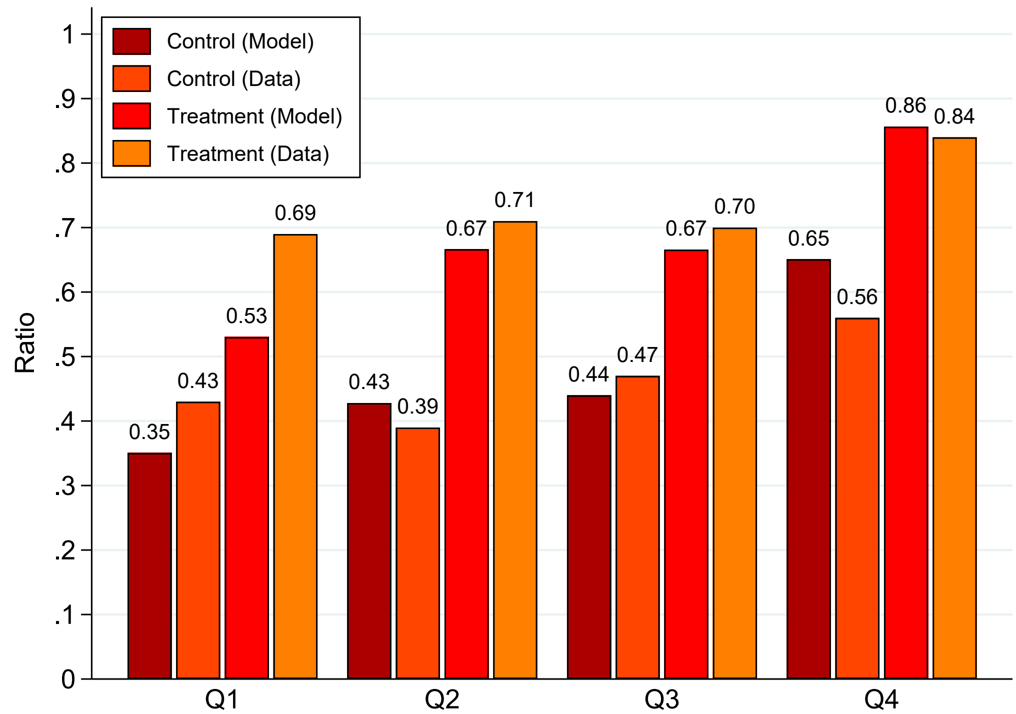

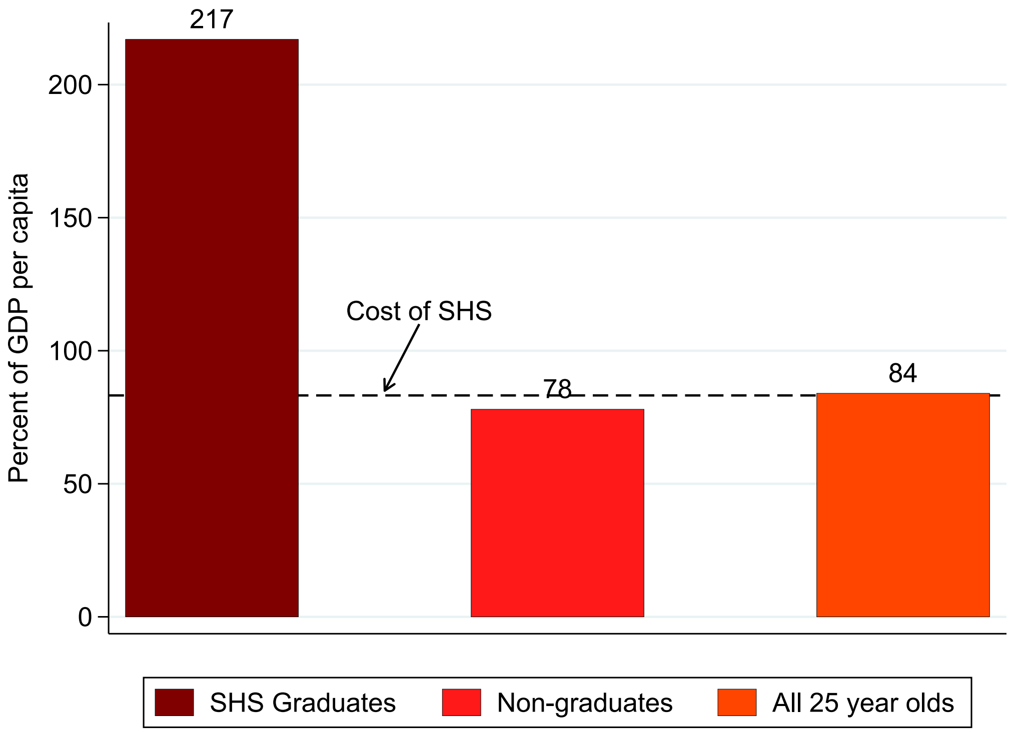

Figure A.2: Compensating Variation of Secondary School to Children

Table A.1: Free Secondary Schooling Policies in Developing Countries

Country	Year	Requirement
Benin	2007	Pass Brevet d'Etudes du Premier Cycle
Gambia	2015	Pass Basic Education Certificate Exam
Ghana	2017	Pass Basic Education Certificate Exam
Kenya	2008	Pass Certificate of Primary Education Exam
Malawi	2019	Pass Primary School Leaving Certificate Exam
Mauritius	2016	Pass General Certificate of Education Exam
Nepal	2018	Pass final district-level exam
Philippines	1988	Do not fail in two consecutive years
Rwanda	2012	Score $\geq$ 'High' on O-level Test
Sierra Leone	2018	Score $\geq$ 6 on Basic Education Certificate Exam
Tanzania	2015	Pass Standard 7 Exam
Uganda	2007	Score $\geq$ 28 in Primary School Leaving Exam
Zambia	2022	Pass Baccalaureate Exam

Note: This table reports the year that each country adopted a free secondary schooling policy and the merit requirement to attend secondary schooling.

Table A.2: Parameter Values Used in Discussion

Description	Parameter	Misallocation Low	Misallocation High
Fertility of Primary School Graduates	ν_J	1	
Fertility of Secondary School Graduates	ν_S	0.2	
Intergenerational Altruism Factor	b	2.5	
Std. Dev. of Exam Score Noise	σ_ε	1	
Gumbel Scale Parameter of Taste Shock	θ	0.5	
Persistence of Ability Shock	ρ	0.8	
Std. Dev. of Ability Shock	σ_υ	0.5	
Std. Dev. of Idiosyncratic Income Shock	σ_ζ	0.3	
Gains from Secondary School	η_S	5.5	14.0
Goods cost of Secondary School	Ψ_S	1.5	5.0
Savings Wedge	χ	0.1	0.2

Note: This table lists the parameter values used for creating the figures in subsection 3.4. The two versions of the model share many parameters and differing only on three key parameters. Parameters not listed here take the value given by Table 1 for both models.

Table A.3: Labor Income Tax Schedule in Ghana

Income	Tax Rates
First 1,008 GHC (=up to 42% of GDP p.c.)	0%
Next 240 GHC (=up to 52% of GDP p.c.)	5%
Next 720 GHC (=up to 82% of GDP p.c.)	10%
Next 14,232 GHC (=up to 675% of GDP p.c.)	17.5%
Exceeding 16,200 GHC ($\geq$ 675% of GDP p.c.)	25%

Note: The table reports the marginal labor tax schedule in Ghana in 2011. It shows, by income in Ghanaian Cedis (GHC), the marginal tax rate assessed on labor income, and the corresponding ratio of GDP per capita in Ghana in 2011.

B. Model Appendix

In this appendix we define the concepts of recursive competitive equilibrium and balanced growth path for our model. Letting X denote the vector of individual state variables $(\tau, a, z_p, s_p, \zeta_p, \delta_J, \delta_S, z_c, s_c, \tilde{z}_c, \zeta_c)$, a recursive competitive equilibrium is defined as follows.

Definition: A **recursive competitive equilibrium** consists of

1. A price system $w_S(f, P)$, $w_U(f, P)$
2. Household value functions $V(X, f, P)$ and policy functions $a'(X, f, P)$, $c(X, f, P)$, $s_c'(X, f, P)$
3. Perceived laws of motion $f' = F(f, P), P' = H(f, P)$

such that

a) V, a', c, s_c' solve the household's optimization problem given w_S, w_U, F, P.

b) For all f, P,

$$w_S(f, P) = (1-\alpha) AK^{\alpha} (N_J)^{\lambda-1} \left[(N_J)^{\lambda} + (N_S)^{\lambda}\right]^{\frac{1-\alpha}{\lambda}-1},$$
$$w_U(f, P) = (1-\alpha) AK^{\alpha} (N_S)^{\lambda-1} \left[(N_J)^{\lambda} + (N_S)^{\lambda}\right]^{\frac{1-\alpha}{\lambda}-1},$$
$$r^* = \alpha AK^{\alpha-1} \left[(N_J)^{\lambda} + (N_S)^{\lambda}\right]^{\frac{1-\alpha}{\lambda}}.$$

c) Markets clear:

$$N_J = \Big[\int_{6\le\tau\le 12, s_p=J} \zeta_p h(z_p, s_p) f(X) dX + \int_{9\le\tau\le 10, s_c'(X,f,P)=J} \zeta_c h(z_c, s_c') f(X) dX \Big] P,$$
$$N_S = \Big[\int_{6\le\tau\le 12, s_p=S} \zeta_p h(z_p, s_p) f(X) dX + \int_{\tau=10, s_c=S} \zeta_c h(z_c, s_c') f(X) dX \Big] P.$$

d) Perceived laws of motion for f and P coincide with those induced from household policy functions a', c, s_c'.

The balanced growth path is a particular type of recursive competitive equilibrium defined below.

Definition: A **balanced growth path** is a recursive competitive equilibrium that satisfies the following properties:

1) Aggregate population grows at a constant rate: $\frac{P'}{P} = \nu$ for some constant $\nu > 0$.
2) The distribution of X is stationary: $f' = f$.
3) The household value and policy functions do not depend on P.

Along the balanced growth path, aggregate population grows but the distribution of households across individual states remains stationary. Further, the household value and policy functions are independent of aggregate population, and thus household behavior remains the same over time conditional on the individual states.

Now we walk through the details of population growth within the model and discuss how model parameters translate to outcomes that are measured in data such as the aggregate population growth rate and the number of children per household. We start with the most general case that applies to any equilibrium whether it satisfies the properties of a balanced growth path or not. Later, we specialize to the case of the balanced growth path to provide more explicit formulas. By definition, the aggregate population growth rate is given by the formula

$$\text{Agg. Pop. Growth Rate} = \frac{\text{\# births} - \text{\# deaths}}{P} \tag{12}$$

Given the aggregate state variables of the economy, f and P, we have the following accounting equations for births and deaths

$$\text{\# births} = \big[\nu_J \int_{s_p=J,\tau=5} f(X)dX + \nu_S \int_{s_p=S,\tau=5} f(X)dX\big]P \tag{13}$$

$$\text{\# deaths} = \big[\int_{\tau=14} f(X)dX\big]P \tag{14}$$

In any given period, the aggregate population growth rate can be computed from state variables as

$$\nu - 1 = \nu_J \int_{s_p=J,\tau=5} f(X)dX + \nu_S \int_{s_p=S,\tau=5} f(X)dX - \int_{\tau=14} f(X)dX \tag{15}$$

Note that as written, $\nu - 1 > 0$ is the aggregate population growth rate such that $P' = \nu P$. To compare to data, it must be converted to an annual percentage growth rate.

Recall that the aggregate population growth rate is constant along the balanced growth path by definition. By leveraging this assumption we can calculate the aggregate population growth rate as a function of educational shares along the balanced growth path analytically. This calculation provides insight into the changes in population dynamics that can be expected due to changes in education. Such changes are important for our general equilibrium analysis.

With the aggregate population growth rate fixed at $\nu - 1$, we know that the ratio of the population of households of age x and households of age y must be given by:

$$\frac{\int_{\tau=x} f(X)dX}{\int_{\tau=y} f(X)dX} = \nu^{y-x} \tag{16}$$

From that fact that $\tau \in \{1, \cdots, 14\}$ and $\int f(X)dX = 1$ because f is a pdf, we can derive that along the balanced growth path with aggregate population growth rate $\nu - 1$ the following equations are true

$$\int_{\tau=14} f(X)dX = \frac{\nu - 1}{\nu^{14} - 1} \tag{17}$$

$$\int_{\tau=5} f(X)dX = \frac{(\nu - 1)\nu^9}{\nu^{14} - 1} \tag{18}$$

Finally, because household policy functions are invariant with respect to P and f is stationary along the balanced growth path, we have that the share of the adult population with a given level of education is the same for all ages. In particular, this implies that the education shares of the parents giving birth this period can be replaced by the aggregate education shares $\hat{J}, \hat{S}$.

$$\hat{J} \equiv \frac{\int_{s_p=J,\tau\geq 5} f(X)dX}{\int_{\tau\geq 5} f(X)dX} = \frac{\int_{s_p=J,\tau=5} f(X)dX}{\int_{\tau=5} f(X)dX} \tag{19}$$

$$\hat{S} \equiv \frac{\int_{s_p=S,\tau\geq 5} f(X)dX}{\int_{\tau\geq 5} f(X)dX} = \frac{\int_{s_p=S,\tau=5} f(X)dX}{\int_{\tau=5} f(X)dX} \tag{20}$$

Combining equations (17) to (20) with equation (15) yields the following equation which describes the aggregate population growth rate along the balanced growth path as an implicit function of the education shares of the population:

$$\nu - 1 = \left[\nu^9\left(\nu_J \hat{J} + \nu_S \hat{S}\right) - 1\right] \frac{\nu - 1}{\nu^{14} - 1} \tag{21}$$

which can be reduced to

$$\nu^5 = \nu_J \hat{J} + \nu_S \hat{S}. \tag{22}$$

One wrinkle not yet addressed is the fact that, as written, the balanced growth path of the model is not an attractor. That is, the model does not necessarily converge over time to the balanced growth path. To see why, consider a simplified model with two generations, each of whom do nothing other than live through their first period of life and, at the end of their second period of life, die and have ν children who become the new first generation. If the initial stocks of age 1 and age 2 agents are N_1 and N_2, the aggregate population growth rate will oscillate between $\frac{(\nu-1)N_2}{N_1+N_2}$ and $\frac{(\nu-1)N_1}{N_1+\nu N_2}$ indefinitely, never converging to a single constant rate, as there is no mechanism to close "gaps" in size between the initial stocks.

To address the computational issues arising from this fact, we assume that a negligibly small fraction of children leave their parents and have their own children one period earlier than the typical timing (that is, at age 20 rather than 25). This slight randomization in timing effectively mixes away any differences in the initial stocks of agents for each generation, ensuring that the model converges to the balanced growth path over time regardless of the initial state. In our computations, we assume the probability that any given child leaves early is 0.1 percent, small enough to ensure that this outcome has minimal impact on parents' decisions.

C. Intuition and Details on Model Identification

An important question is which of the targeted moments are most informative for each of the estimated parameter values. To help answer this question, we follow Kaboski and Townsend (2011) and compute the percent change in each moment when each parameter is increased by one percent. We then re-estimate the model multiple times under alternative values of the targeted moments, to see how sensitive each parameter is to each target, following Andrews et al. (2017). While in general all moments jointly discipline all the parameters, some parameters correspond more closely to certain moments. The Jacobian matrix is presented in Appendix Table C.1, and the sensitivity matrix is presented in Appendix Table C.2; here, we summarize what we see as the main lessons from these matrices.

The population growth parameters ν_J and ν_S are, perhaps unsurprisingly, significant determinants of the aggregate population growth rate and the treatment effects on fertility. The variance and persistence parameters of the ability process, σ_v and ρ, naturally increase the variance of the permanent component of income and the intergenerational schooling correlation, but also have sizable effects on many other moments in equilibrium.

The effectiveness of schooling, η_S, and the intergenerational altruism parameter, b, govern the benefits of schooling and thus result in similar changes, notably a sizable increase in aggregate secondary attendance. The key difference is that η_S increases the treatment effect on human capital while b has a minimal effect, as it only impacts the parent's valuation of better schooling. Intuitively, the cost of schooling, Ψ_S, decreases school attendance, increases the treatment effect on schooling, and consequently increases (in absolute value) the treatment effect on fertility.

Finally, the savings wedge, χ, and the standard deviations of the test score noise and taste shocks, σ_ε and θ, all jointly impact secondary attendance in the top and bottom quartiles of the test score distribution as well as the difference in treatment effect between the quartiles. In fact this was the purpose of introducing these shocks into the model, and without them schooling completion and the treatment effect on schooling are always (counterfactually) much larger for those with higher test scores.

Table C.1: Elasticities of Moments to Parameters

	ν_J	ν_S	σ_υ	ρ	η_S	b	Ψ_S	σ_ε	θ	χ
Aggregate population growth	0.8	0.0	0.0	0.6	-0.2	-0.1	0.0	0.0	0.0	0.0
Aggregate SHS attendance	-0.5	0.1	-0.1	-2.2	0.9	0.3	-0.1	-0.1	0.0	0.0
Intergenerational school corr.	0.0	0.0	0.5	1.7	-0.7	-0.2	0.2	-0.4	0.0	0.0
Var(permanent income)	-0.4	0.1	0.8	-1.1	1.7	0.0	0.0	0.0	0.0	0.0
SHS in top quartile	0.0	0.0	0.2	-0.4	0.3	0.0	0.0	-0.2	0.0	-0.1
SHS in bot quartile	0.0	-0.2	-0.4	-1.0	0.6	0.3	0.0	0.2	0.0	0.0
TE on human capital	0.0	0.0	-0.1	-0.2	0.1	0.0	0.0	0.0	0.0	0.0
TE on fertility	-0.5	0.3	0.8	3.0	-1.9	0.1	-0.4	0.5	0.0	-0.1
TE on SHS completion	-0.4	0.0	-0.8	-2.8	1.8	-0.2	0.4	-0.5	0.0	0.1
TE on SHS, Q4-Q1 difference	1.0	-0.3	-0.5	12.8	-8.8	2.7	-1.5	3.6	-1.6	1.1

Note: This matrix represents the elasticities of each moment to each parameter. The entry in row r and column c represents the percentage change in model moment r resulting from a one-percent increase in model parameter c.

Table C.2: Sensitivity of Parameters to Moments

	ν_J	ν_S	σ_υ	ρ	η_S	b	Ψ_S	σ_ε	θ	χ
(i) TE on human capital x 1.5	3.6	2.1	-1.0	0.0	-0.7	4.2	-1.7	2.8	1.3	3.0
(ii) TE on fertility x 1.5	8.4	-3.1	4.5	-0.4	6.9	-2.9	4.7	3.8	4.7	-2.2
(iii) TE on SHS completion x 1.5	6.4	3.1	-3.0	-1.1	3.3	0.2	1.2	-0.7	1.4	8.3
(iv) TE on SHS, Q4-Q1 difference x -1	2.3	1.8	1.0	0.2	-0.6	-1.9	7.1	2.2	-2.8	-7.6
(i), (ii), (iii) and (iv) combined	10.3	-5.9	-2.8	1.5	17.9	-4.8	28.9	7.0	-2.8	-4.4

Note: This matrix displays the percent change in each parameter when the mode is re-estimated to match different target moments. Row (i) is when we target a fifty percent higher target for the treatment effect on human capital but keep all other targets the same. Row (ii) is when we target a fifty percent higher treatment effect on fertility. Row (iii) is when we target a treatment effect on secondary school completion that is fifty percent higher. Row (iv) is when we target a difference in treatment effects on schooling between Q4 and Q1 that is -4 instead of 4. The bottom row is when we re-estimate the model to match all of the higher targets..

C.1. Robustness

In a similar vein as the above exercises, here we confirm that the fact the estimated model undershoots the intergenerational schooling correlation and the treatment effect on schooling attendance does not substantially alter the model's conclusions.

First, we increase the cost of schooling parameter Ψ_S by 40 percent which allows the model to match the treatment effect on schooling, albeit at the expense of a slightly worse fit to other model moments. Under this alternative parameterization, the free schooling policy results in a 0.8 percent decline in GDP (vs 1.0 percent in the baseline) and a 14 percent increase in school attendance (vs 12 percent baseline). Neither of these results are substantially different from the baseline.

We also perform a similar exercise by increasing the parameter governing the persistence of the ability ρ process by 15 percent to match the intergenerational schooling correlation coefficient (again with a slightly worse fit for other moments). In this case, the free schooling policy results in a 2.0 percent decline in GDP and only a 4 percent increase in attendance. The intutition for this result is that a higher persistence parameter implies less misallocation as fewer high ability children are born to poor families. Thus missing on this moment means, if anything, that our conclusions are too generous to the free schooling policy.