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

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

July 2025

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 leads to only a modest gain in GDP per capita. The human capital gains from expanded education access are offset by lost income during schooling years and dampened by negative selection of new students entering secondary school. An alternative policy that spends the same resources improving school quality has significantly larger effects.

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

Across the developing world, millions of children never advance beyond primary school. Low secondary enrollment rates are most commonly attributed to credit constraints that prevent poorer households from investing in education. To the extent that this is true, low secondary enrollments suggest a significant misallocation of talent as bright children are forced to cut their education short (Hsieh, Hurst, Jones, and Klenow, 2019). Motivated by this concern, many developing nations are implementing publicly funded ('free') secondary schooling programs designed to boost enrollment by removing financial barriers for poor families (Center for Global Development, 2022).

To help evaluate whether free secondary schooling delivers on its promise, we build and estimate a macroeconomic model that is tailored to capture key realities of the educational landscape in developing countries. Our analysis builds on established frameworks for analyzing credit constraints and human capital investments (Galor and Zeira, 1993; Bénabou, 2002) and adds three additional features. First, attending secondary school means forgoing significant labor income—a major consideration given how many teenagers in Africa and South Asia work to support their families. Second, school access depends not just on financial resources but also on passing an entrance exam, creating an additional selection mechanism. Third, education not only raises individual labor earnings but also alters family sizes in the subsequent generation, through changes in fertility and child mortality rates. Incorporating these elements allows us to better trace through the long-run general equilibrium consequences of removing financial barriers to secondary education.

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, 2025). Those offered the scholarships were about 25 percent more likely to finish secondary school than a control group four years later. The scholarships raised literacy and mathematics test scores by about 0.2 standard deviations, which is comparable to the effects found in other successful education interventions. A decade after graduation, child mortality among the children of scholarship winners fell by a substantial 45 percent, increasing net fertility despite no change in birth rates (Duflo, Dupas, Spelke, and Walsh, 2024). This intervention, and the dedicated efforts to follow up with participants for nearly two decades, represent the longest-running empirical evaluation of secondary school scholarships in existence.

We use our estimated model to simulate the long-run aggregate effects of permanently enacting free secondary schooling financed through progressive taxation. The model predicts a one-fourth increase in secondary school graduates, broadly consistent with the substantial enrollment gains described by the Center for Global Development (2022). However, these educational gains translate into surprisingly modest economic effects: GDP per capita increases by only about 2 percent in the long run. Meanwhile, financing the policy requires tax increases of around 3 percent of GDP, meaning the policy generates benefits worth only two-thirds of its cost. The result is lower average consumption levels — a finding inconsistent with free schooling being a force to raise living standards by eliminating talent misallocation.

To understand why the model makes such pessimistic predictions, we conduct a series of counterfactual experiments using the estimated model. We find that a substantial portion of the earnings gains from free school are offset by lost earnings during secondary school years in our main analysis: when we eliminate forgone earnings by allowing students to work while attending school, the GDP gains from free schooling nearly double to 4 percent. Negative selection of new students plays an even larger role. The RCT targeted poor but high-ability students, but free schooling would predominantly attract students of more marginal ability. When we assume all new attendees match the RCT participants' skill distribution, the GDP gains jump to 6 percent — consistent with Hendricks and Schoellman (2014)'s findings of significant ability sorting in U.S. higher education. Demographic effects provide a modest offset: eliminating the net fertility differences by education level reduces the GDP gains of free schooling by 0.3 percentage points (pp), as fewer high-ability children are born to educated parents.

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 10 percent in the long run. Taxes per capita would still increase by about 3 percent of GDP, meaning that the free schooling policy would easily pay for itself in the long run and raise average consumption levels. The sizable GDP per capita gains in this counterfactual illustrate how the model's rather negative 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 and negative selection of new entrants to school.

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 around 10 percent per capita, around five times as large as free schooling, and comfortably enough to cover its costs in the long run.

We conclude that improving the quality of secondary schools is a more effective way to raise living standards than making the current schools free and that low secondary schooling levels are to a large extent an efficient response to low quality school options. Though experimental evaluations find substantial impacts from free schooling, it is also fairly costly (particularly when compared to quality-improving interventions) and largely influences marginal students, rather than those who would experience large earnings gains. These implications are 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 contributes to the large literature focused on human capital accumulation in macroeconomics.[1] Our paper is arguably most closely related to the studies of Abbott, Gallipoli, Meghir, and Violante (2019), Daruich (2020), and Krueger, Ludwig, and Popova (2025) who reach fairly positive conclusions about the effects of expanding educational access in the United States. Our more negative conclusions likely stem from the lower quality of schools in Africa; improving school quality also comes out ahead of free schooling in the analysis of Krueger et al. (2025). Our paper complements the empirical work by Khanna (2023), who estimates the impact of expanding education including the negative wage effects of increasing the relative supply of skilled labor. As in Krueger and Lindahl

[1]See e.g. Becker and Tomes (1979); Loury (1981); Galor and Zeira (1993); Bils and Klenow (2000); Restuccia and Urrutia (2004); Erosa, Koreshkova, and Restuccia (2010); Lochner and Monge-Naranjo (2011); Manuelli and Seshadri (2014); Hendricks and Schoellman (2018); Lee and Seshadri (2019); Porzio, Rossi, and Santangelo (2022); Buera, Kaboski, Rogerson, and Vizcaino (2022).

(2001) and Pritchett (2001), our study attempts to learn about macroeconomic effects of educational expansions with microeconomic estimates of the private returns to education.[2]

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

This section examines both macro and micro evidence on secondary schooling outcomes and free education policies in developing nations. The patterns we highlight directly shape our modeling framework and the counterfactual scenarios we explore in subsequent sections.

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.

[2] Our paper also builds on the recent literature attempting to quantify the extent to which credit market imperfections drive misallocation in developing countries. To cite just a few, 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.

Figure 1: Primary and Secondary School Enrollment Rates

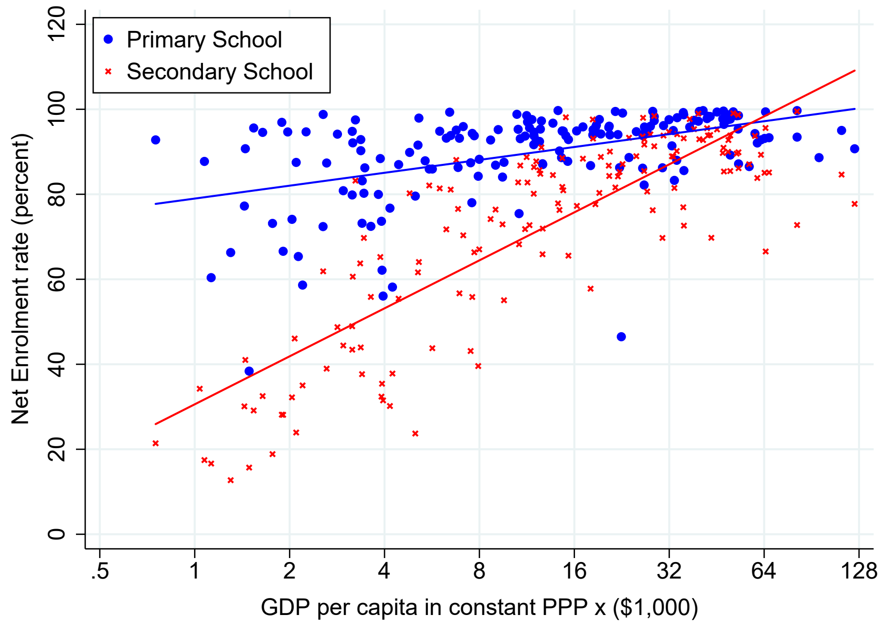

One salient difference between rich and poor countries in terms of education policy is that richer countries are much more likely to publicly finance secondary education. It is not surprising, then, that many developing countries have recently considered implementing 'free' schooling policies, in which 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 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 conflict 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 identify credit constraints as a primary barrier to enrollment.

Duflo et al. (2025) conducted the first long-run randomized 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 also exhibited higher human capital as measured by math and reading tests. The magnitudes were substantial, and translated into 7.6 percent higher human capital per year of school completed. The authors followed up with the study participants regularly for more than a decade after graduation. At that point, many of the participants had their own children, allowing the authors to estimate the effects of secondary school scholarships on child health and development outcomes. Rates of fertility were not statistically significantly different between the two groups, but the children of scholarship winners had 45 percent lower overall mortality rates than the control group, highlighting the impact of secondary education on net fertility (Duflo et al., 2024). 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 6.4 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 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 à 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 costs of schooling, in the form of lost labor income and fertility rates that depend on education levels. As in most models in this literature, ability is passed down imperfectly from parent to child to allow for misallocation of talent through borrowing constraints that keep high-ability children born to poor parents out of school. As in a growing literature in macro development, we allow for saving constraints that impede households from simply saving themselves out of borrowing constraints (see e.g. Donovan, 2021), allowing this misallocation to be potentially large.

As we will demonstrate later in Section 3.4, including experimental moments as targets in model estimation is very helpful in disciplining quantitatively the extent of misallocation. To leverage this advantage, we include several model features that make it possible to replicate and match a particular experimental evaluation of free schooling. Attending secondary school requires passing an admissions exam with a student's score depending on their ability, allowing us to replicate experimental selection criteria and moments based on exam scores. Further, schooling decisions are subject to taste shocks, proxying for unobserved factors that may lead high-ability children not to attend school even if it were offered to them for free, allowing us to replicate non-perfect compliance, as observed in reality.

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 five 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 when schooling is made free. In the fifth period, all children work at their respective education level, which is fixed afterward.

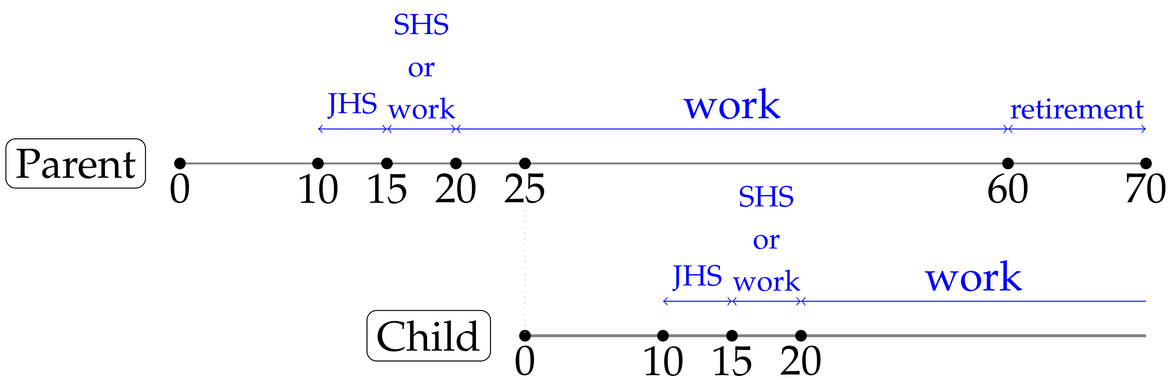

At the beginning of period six, when turning age 25, children leave their parents, have children, and become parents themselves. We abstract from household formation decisions since they do not seem crucial for our task at hand. Instead, we model households as continuous dynasties that do not mix. These 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 $1 + \nu_{s_p}$, changes with the parent's schooling level s_p. As a result, policies that increase an individual's level of schooling will alter their fertility, consistent with a variety of evidence for developing countries and eventually, the experimental evidence that we use in model estimation.

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 $N(0, \sigma_\upsilon^2)$. 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 $U(c) = \log(c)$ 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).

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 $N(0, \sigma_\varepsilon^2)$.

Parents make schooling decisions for their children when the children turn age 15, after observing the children's ability and test scores. We also allow for 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. These shocks serve as proxy for all unobserved, unmodeled features that may lead otherwise observably identical parents to make different schooling decisions. These shocks end up being necessary to match observed experimental outcomes, as the model would otherwise predict counterfactually high correlations between test scores, attendance, and treatment effects.

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.

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.

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.

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.

The labor income y of an individual equals the product of four 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). Finally, $\omega(\tau)$ represents an age-dependent term, giving rise to a deterministic life-cycle earnings profile; such a term is necessary to capture the idea that school-aged children may earn less than adults, influencing the opportunity cost of schooling. For example, the labor income of an individual with age $\tau = 9$ and education level S is given by

$$y(z, S, \zeta, 9) = w^S \zeta h(z, S)\, \omega(9) = w^S \zeta z \eta_S \omega(9). \tag{5}$$

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

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

where the maximization is subject to the flow budget constraint

$$a' + c + (1 + \nu_{s_p}) \mathbf{I}(s'_c = S)\, \Psi_S = y_p(z_p, s_p, \zeta_p, 9) + (1 + r)a + (1 + \nu_{s_p})(1 - \mathbf{I}(s'_c = S)) y_c(z_c, J, \zeta_c, 1) - T(z_p, s_p, \zeta_p, z_c, J, s'_c, \zeta_c) \tag{6}$$

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 $(z_p, s_p, \zeta_p, z_c, J, s'_c, \zeta_c)$.[3] We suppress the dependence of y_p, y_c, and T on f and

[3]Note that T depends on both the children's current schooling level s_c $(= J)$ and next period's

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 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

$$a' + c = y_p(z_p, s_p, \zeta_p, 10) + (1+r)a + \left(1+\nu_{s_p}\right) y_c(z_c, s_c, \zeta_c, 2) - T(z_p, s_p, \zeta_p, z_c, s_c, s_c', \zeta_c), \tag{8}$$

$f' = F(f, P)$, $P' = H(f, P)$, 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.

schooling level s_c'. This is because the labor income depends on educational attainment, and only the children who do not go to school ($s_c' = s_c$) earn labor income in the current period.

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, dubbed the *high misallocation economy*, has a higher cost of school, higher schooling quality, and a larger savings wedge. [4]

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 misalloca-

[4] In particular, the low misallocation economy features values $(\Psi_S, \eta_S, \chi, \sigma_v)$ of $(3.5, 1.5, 0.1, 0.6)$ while the high misallocation economy has values $(10.0, 4.0, 0.2, 0.25)$. The remaining parameters, which are not crucial for the conclusions in this section, are set to their estimated values from Section 4 below.

Figure 2: Child SHS Attendance Probability

Low Misallocation

High Misallocation

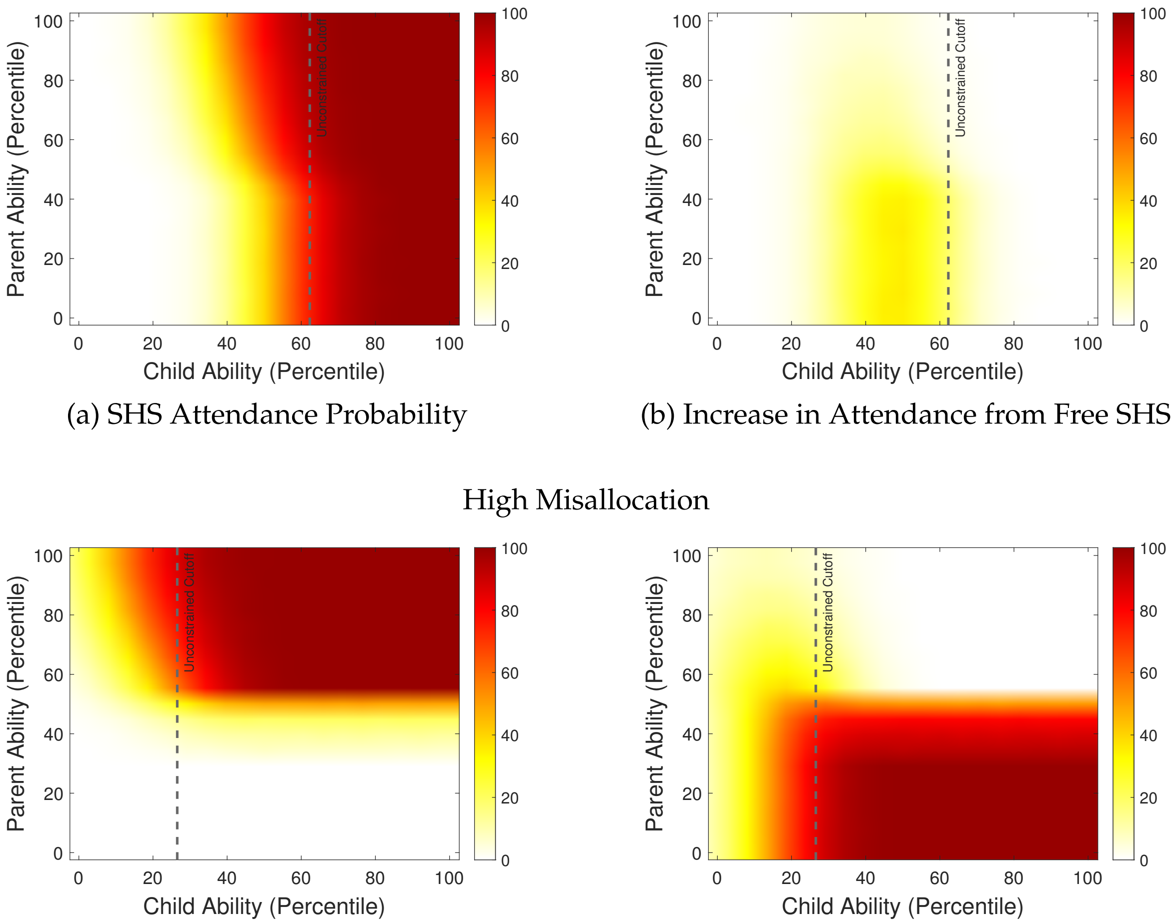

(a) SHS Attendance Probability

(b) Increase in Attendance from Free SHS

(c) SHS Attendance Probability

(d) Increase in Attendance from Free SHS

tion 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 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.

It is difficult, however, to distinguish between the two economies using aggregate data alone — they exhibit essentially identical secondary school attendance rates and earnings premiums (Appendix Table A.2 provides the exact values along with other statistics from these economies). In the low misallocation economy, attendance is low because the return to schooling is low, but those who do opt to attend are high ability so that the observed wage premium remains elevated. In contrast, the high misallocation economy exhibits higher quality schooling, but attendance is low because many of these children face borrowing constraints. Attendance is largely determined by parental income, rather than child ability, dragging down the earnings premium. This motivates our use of experimental moments in model estimation, which allows us to overcome the unobservability of student ability and effectively distinguishes between the two cases. Panels (b) and (d) show the increase in the probability of attendance when SHS is offered for free. In the low misallocation economy, the increase in attendance is small and concentrated among marginal children; the result is modest 11 pp increase in attendance and 2.4 percent increase in earnings. In contrast, these increase by 50 pp and 28 percent respectively under high levels of misallocation where there are large increases in attendance even at the top of the ability distribution. We thus view incorporating experimental results as crucial in determining the scope for earnings gains from free schooling policies.

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.

We normalize the age-dependent component of income, $\omega(\tau)$, to unity for all age periods other than the first. This is consistent with the general finding that life-cycle earnings profiles in developing countries are largely flat (Lagakos, Moll, Porzio, Qian, and Schoellman, 2018), and the short two-period panel provided by the control group from Duflo et al. (2025) (see Appendix Table A.3). We choose the initial productivity level, $\omega(1)$, to be 0.66 following the age-employment-rate profiles in Ghana's Living Standard Survey (Appendix Table A.4). Employment rates for adults are high at all ages except for the youngest group (ages 15-20) who, conditional on not being in school, have an average employment rate that is two-thirds that of other adults.

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.5, 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.

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:

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

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.

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 aggregate secondary school completion rate for all school-aged individuals (taken from World Bank Education Statistics for 2014 which, of the available years, most closely responds to the graduation year of 2011 in the experimental sample) and the secondary school completion rates in the top and bottom test score quartiles of the control group of Duflo et al. (2025). The final non-experimental moment we target is the variance of the permanent component of log wages.[5] 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 et al. (2025), 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 fol-

[5] 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 $Cov\left(\log y_{i,t}, \log y_{i,t-1}\right) = Var\left(\log\left(w_{i,t}h_i\right)\right)$.

lowing 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" who passed the BECE (the entrance exam) but did not end up enrolled in secondary high school the following year (95 percent cite financial difficulties as the reason for non-enrollment). This is important, as it leads the experimental sample to be positively selected on ability and, consequently, returns to schooling. Appendix Table A.6 demonstrates this by comparing eventual outcomes (at age $\sim$28) in the experimental control group to an equivalent (age 25-30) representative sample from the Ghanaian Living Standards Survey (GLSS). The experimental control group has higher average years of education (11.4 vs 7.6), is more like to have completed secondary school (45 percent vs 21 percent), is less likely to be self-employed (25 percent vs 36 percent), and is more like to have a wage job (35 percent vs 26 percent). If we do not account for this and incorrectly interpret the results of the experiment as applying to the average student, we will overstate the gains from schooling. Thus this channel of "selection" is one reason that extrapolating directly from these experimental results will overstate the aggregate gains of the policy, a fact which we return to below.

We choose a passing test score cutoff so that only the top 42 percent of students in the model pass, consistent with the actual BECE passing rate. It is more difficult to literally match the requirement that the student did not enroll in the following year; many of these students did eventually complete SHS, even in the control group, so selecting model students who passed but do not attend would dramatically understate control group completion. Instead, since the vast majority of selected students do not enroll for financial reasons, our strategy is to choose a parental income cutoff so that the control group completion rate in the model matches that observed in the experiment (45 percent). We then choose the experimental sample in our model to be a subset of those with test scores *above* the test score cutoff and parental 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 treatment effects on net fertility and human capital, which we view as the most important – and precisely estimated – findings of the experiment. The impact on net fertility (which is discussed at length in Duflo et al., 2024) is positive. While there is no long-run effect on the number of children born, treated women experience a 45 percent reduction in child mortality. Together with the control rates of child mortality (4 percent for women, 2.9 percent for men) and no apparent impact on treated men, this implies an average net fertility increase of 0.9 pp, which is our target. In terms of human capital, 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. (2025) 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 estimated impact is 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). 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 7.4 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 pp. 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 pp, 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 Parameter Values

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 quite good. The model does well in matching the treatment effects on human capital (8.6 percent versus 7.4 percent in the data) and fertility (0.9 percent versus 0.9 percent in the data). We do miss slightly on the variation between test scores and treatment effects with the model generating a treatment effect difference between score quartiles of about zero (i.e. flat) compared to the 4 pp targeted. Still, the model falls well within the 95 percent confidence interval for the moment. Appendix Figure A.1 displays SHS completion rates for all test score quartiles (including untargeted quartiles) demonstrating that, broadly speaking, the model gets these patterns correct.

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

Table 2: Targeted Moments and Model Predictions

Moments	Data	Model
Aggregate Population Growth	2.2	2.2
Aggregate SHS Completion Rate	35	34
SHS Completion, Q4 of Test (Control Group)	56	58
SHS Completion, Q1 of Test (Control Group)	43	40
Var(Permanent Component of Income)	0.22	0.25
Treatment Effect on Human Capital	7.4 (3.7, 11)	8.6
Treatment Effect on Fertility	0.9 (0.0, 1.7)	0.9
Treatment Effect on SHS Completion	27 (22.0, 30.8)	27
Treatment Effect on SHS Completion, Q4 - Q1	4 (−7.2, 13.7)	0
Intergenerational Schooling Correlation	0.45 (0.43, 0.46)	0.44

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.

that the model is precisely estimated in a statistical sense.

The estimated parameters seem reasonable from an economic sense as well. The estimated fertility parameters imply that each family has around 2(1+0.75) = 3.5 children which is roughly in line with the figure from the 2008 Demographics and Health Survey for Ghana for 4.0. We do note, however, that in order to match the positive observed treatment effect on fertility, the model assigns a larger number of children to those with secondary education than primary education ($\nu_S > \nu_J$). This contradicts the aggregate birthrates observed in the DHS, which show average fertility of 4.1 children for women with JHS only, and 2.6 children for those with secondary education or more. Still, we do not view this as a problem for the purposes of our paper where only *shifts* in birthrates matter for policy counterfactuals.

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 6.3 percent per year for this experimental sample. This is generally in line with other estimates of returns

Table 3: Parameter Estimates and Confidence Intervals

Parameter	Description	Estimate (Confidence Interval)	Closest Moment
ν_J	Fert. of Primary Graduates	0.71 $(0.68,\ 0.72)$	Agg. Pop. Growth
ν_S	Fert. of Secondary Graduates	0.77 $(0.72,\ 0.79)$	TE on Fert.
ρ	Persistence of Ability Process	0.82 $(0.77,\ 0.88)$	Intergen. Corr.
σ_υ	Std. Deviation of Ability Shock	0.32 $(0.24,\ 0.36)$	Var(Perm. Income)
Ψ_S	Goods Cost of Secondary School	2.17 $(1.96,\ 2.48)$	TE on SHS Compl.
η_S	Efficiency of Secondary School	4.67 $(3.97,\ 6.17)$	TE on HC
b	Intergenerational Altruism	1.30 $(1.2,\ 1.7)$	Agg. SHS Compl.
σ_ε	Std. Deviation of Exam Score	1.29 $(1.15,\ 1.51)$	TE Compl. Q4-Q1
θ	Gumbel Parameter of Taste Shock	0.24 $(0.20,\ 0.30)$	SHS Compl. Q1
χ	Savings Wedge	0.13 $(0.10,\ 0.13)$	SHS Compl. Q4

Note: This table reports the estimated parameters. The confidence interval is the 2.5th and 97.5th percentiles of 100 bootstrapped parameter estimates as well as the target moment most closely corresponding to each parameter.

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 35 percent of GDP per capita, which is substantial, though somewhat lower than 70 percent reported by Duflo et al. (2025)[6]

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 compensating variation is substantially larger than the cost of schooling for children who attend SHS and

[6]The fact that tuition is, in reality, collected yearly over the course of schooling and rather than entirely up-front as in the model (i.e. due to the fact that each period is 5 years) may explain our lower value, particularly in the presence of borrowing constraints where large up-front payments are very costly.

modestly lower for children who do not. 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 0.13, which implies that households save at around 8.3 percent per period, or 1.6 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).[7]

The estimated value for the intergenerational persistence of ability, ρ, is 0.82, 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_υ, is 0.32. While not directly interpretable, this value (along with the other parameters) generates a Gini coefficient within the model of 0.32. 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.

Parameter Identification: To generate some initial intuition on the role that each moment plays in disciplining parameters, Appendix Table C.1 displays the elasticity of each model moment to each parameter. We summarized the lessons from this matrix here and in the column of Table 3 labeled "Closest Moments". Unsurprisingly, the parameters governing the number of children born based on education level, ν_J and ν_S largely impact (and thus are identified by) the aggregate population growth rate and the experimental treatment effect on fertility. Together, the intergenerational schooling correlation and the variance in permanent income levels determine the persistence and variance parameters of the AR(1) process governing intergenerational ability transmission, ρ and σ_υ.

Conditional on the previous four moments, the two most important parameters for the model, the cost Ψ_S and effectiveness η_S of schooling are jointly determined by the experimental treatment effects on school completion and human capital. Despite being jointly determined, we prefer to think of the cost Ψ_S as being determined by the effect on attendance (as a higher cost means free schooling represents a larger

[7]The average household asset-to-income ratio in the model is 0.7. 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.

shift) and effectiveness η_S as being driven by the effect on human capital, conditional on attendance. Once the costs and benefits of schooling are pinned down, the intergenerational altruism parameter b determines the aggregate level of SHS attendance. The final moments, the degree of noise in test scores σ_ε, the Gumbel parameter for the taste shock θ, and the savings wedge χ are determined by the moments describing how completion and treatment effects vary at the bottom and top of the test score distribution.

5. Simulating the Effects of Free Secondary School

Using the estimated model, we simulate the effects of a national free secondary schooling policy that reduces the cost of secondary schooling faced by households, Ψ_S to zero with these costs instead paid by the government. 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.5).

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.[8]

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 ex-

[8] 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.

perimental moments in particular as these are estimated with large standard errors.

The number of secondary schooling graduates increases by about 8 pp under free schooling, from 34 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 rises due to the schooling expansion, and the population growth rate increases very slightly by 0.01 pp. Adult earnings increase by about 2.5 percent from the policy, stemming largely from the higher wages for the 8 percent of the population now receiving secondary education. Child earnings increase by less, only 0.6 percent, reflecting the opportunity cost of higher schooling — earnings among specifically SHS-aged children fall by 9 percent.

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

Statistic	Change Under Free Schooling
GDP per Capita (%)	2.1 (1.0, 3.1)
Secondary School Completion (pp)	8.0 (6.3, 10.9)
Population Growth Rate (pp)	0.01 (-0.00, 0.02)
Adult Earnings (%)	2.5 (1.5, 4.0)
Child Earnings (%)	0.6 (-1.2, 0.9)
Taxes per Capita (%)	3.0 (2.7, 3.6)
Skilled Wage/Unskilled Wage (%)	-7.3 (-10.5, -2.7)
Gini Coefficient	-0.02 (-0.03, -0.01)

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.

GDP per capita increases by only about 2.1 percent in the long-run, less than one-third of the 6.4 percent increase implied by the back-of-the-envelope calculations of Section 2. Our confidence interval excludes gains larger than 3.1 percent. The long-run cost of the policy is 3.0 percent of GDP, larger than the GDP gains, implying aggregate consumption declines. Relative wages of the skilled fall by about 7 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.[9]

Figure 3: Child SHS Attendance Probability in Estimated Model

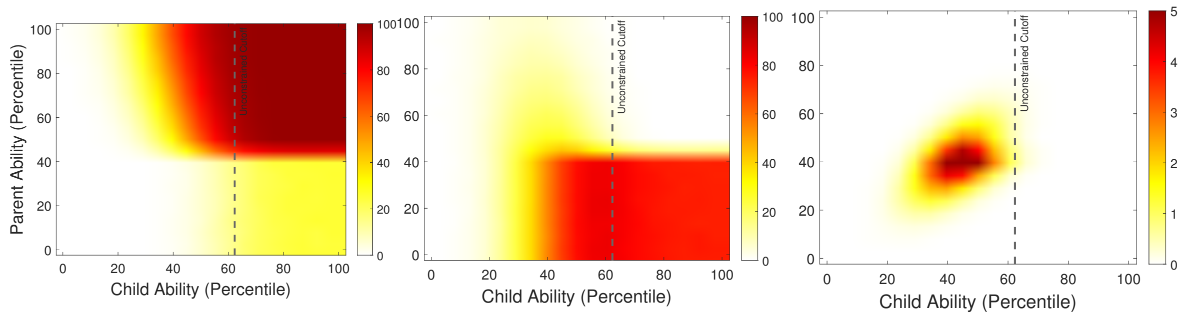

(a) Attendance Prob.

(b) ↑ in Attendance

(c) Pop.-Weighted ↑ in Attend.

The small-to-modest increase in GDP suggest that the estimated model does not feature particularly 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 panels (a) and (b) display identical plots using the fully estimated model which, at a glance, appears to fall between the two extremes shown in Figure 2. While attendance does have some dependence on ability, many of the highest ability children attend SHS even if they are born to low ability parents. Like the high-misallocation economy, the increase in schooling (panel b) is substantial for high-ability children born to low-ability/poor parents (though the magnitude is smaller), but like the low-misallocation economy, the optimal unconstrained attendance threshold is quite high, with free schooling inducing many who are below

[9]Appendix D discusses the impact of the policy on welfare. This discussion is quite involved as the presence of endogenous fertility inherently makes the idea of welfare difficult; thus we omit it from the main text.

the threshold to attend. Further, the estimated correlation between parent and child abilities is quite high. Unlike panel (b), which shows the increase in attendance probability for any hypothetical parent-child ability combination, panel (c) displays the distribution of those who are induced to attend by the policy[10], accounting for the underlying distribution of abilities and reflecting the actual composition of new attendees. Here we see that higher attendance is almost entirely concentrated below the cutoff. Perhaps the most accurate description of this economy, then, is one in which the *potential* for misallocation is high in that a hypothetical high-ability child born to poor parents will be kept out of school, but because of low estimated schooling quality (high cutoff) and high estimated intergenerational ability persistence (low population in lower-right quadrant), equilibrium misallocation is low.

Comparison to Actual Free SHS Rollout: While our model is estimated to match Ghana circa 2008, when secondary school was costly, Ghana eventually implemented a universal free SHS policy in 2017 providing an opportunity to evaluate the model's quantitative predictions. The main difficulty in making this comparison is that attendance was increasingly rapidly even before the policy was implemented and, as it was universal across the country, there are no natural control groups to serve as a counterfactual in order to estimate the policy's impact. We instead compare the growth in schooling attendance to three different hypothetical counterfactuals — the best-fit trendline for attendance from 2011 to 2017, the trendline from 2013 to 2017, and finally the assumption that there is no trend and attendance growth from 2016 to 2017 would have been zero absent the policy (i.e. all 2016-to-2017 growth is attributable to the policy). These trendlines, plotted against actual enrollment data from 2011 to 2017, are plotted in Appendix Figure A.2.

Depending on which counterfactual is chosen, the increase in enrollment due to the free schooling policy ranges from 2.6 percent (trendline since 2011) to 17.7 percent (no trendline). The full results are displayed in Appendix Table A.7. This is a wide range, but the model's predicted increase of 8 percent is consistent with these estimates.

5.2. Opportunity Cost, Selection, and Differential Fertility

Given that the back-of-the-envelope calculation at the end of Section 2 suggested a GDP gain of around 6 percent based on the experimental estimates, it is natural to wonder what mechanisms lead the model to predict a substantially smaller 2.1 percent gain. Here we perform counterfactual experiments using the model to

[10]The units are normalized to to sum to 100, representing 100 percent of new attendees.

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.

We focus on the channels of opportunity cost, selection (i.e. the fact that the experimental sample is highly selected on ability relative to the population at large), and fertility. For each channel, we conduct a counterfactual exercise where we shut down the channel's impact for new SHS attendees (i.e. changes are made only for those induced to attend SHS by the free schooling policy, keeping unaffected those who would have attended regardless of the policy) and reevaluate the impact of the free schooling policy.[11] More precisely, we start from the pre-expansion steady state and implement the free secondary school policy as in the baseline; however, each child who would not have gone to school according to the pre-expansion policy functions has their schooling decision changed in some way, depending on the channel being examined. For opportunity cost, we eliminate the tradeoff between education and work and allow new SHS attendees to continue to work at full productivity while at school, effectively eliminating the opportunity cost of attendance. For selection, we take the ability distribution of the experimental sample and have (potential) new attendees redraw their ability z from this distribution before making their attendance decision; the result is that new attendees now exhibit the same earnings gains as the positively-selected experimental sample, on average. For fertility, we eliminate any changes in the number of children for new attendees (i.e. new attendees maintain fertility ν_J even after attending), effectively eliminating the impact of the policy on fertility.

Beginning with opportunity cost, the results from the first counterfactual where new attendees can continue to work at full productivity while attending school can be seen in the dark orange columns of Figure 4, marked 'no opportunity cost.' In this counterfactual, GDP per capita increases by 3.8 percent (compared to the 2.1 percent

[11]Shutting down the channel only for new attendees implies that counterfactual exercise begins from the same pre-policy balanced growth path as the baseline model. This eases comparisons by ensuring that any differences in policy impact arise from changes in the results of the policy rather than changes in the pre-policy baseline.

in the baseline). This larger income gain is driven in part by substantially larger increases in secondary attendance of 12.8 pp (vs 8.0 pp baseline). This larger increase is entirely due to a substantial increase in children's earnings which now rise by 16.2 percent (vs 0.6 percent) while the gains to parents' earnings are essentially unchanged. Overall, this counterfactual suggests that opportunity cost reduces the gains from free schooling by 1.7 pp (3.8 percent minus 2.1 percent). We can further decompose this 1.7 pp in order to understand how much of this difference arises from the fact that eliminating opportunity cost mechanically increases the GDP impact, as childrens' lost income no longer drags gains down, and how much occurs endogenously from higher education levels (as parents are now more likely to send their child to school). Performing this decomposition shows that the entirety of the gains (1.7 pp) can be attributed to the mechanical, rather than endogenous, channel; the students induced into attendance by eliminating opportunity cost are entirely marginal and experience no net income gains, further highlighting the conclusion above that it is largely low school quality that explains low attendance rates.

The results of the second counterfactual, where the abilities of new attendees are modified to match the distribution observed in the experimental sample, are displayed as the red bars marked 'No selection on ability' in Figure 4. Here the gains to GDP climb substantially from 2.1 to 6.0 percent and the increase in SHS attendance almost doubles to 15.4 percent. This effect arises due to the fact that the experimental sample was positively selected on ability relative to the average potential new attendee in the baseline counterfactual. When the average ability of new attendees is increased to match the experimental sample, the gains from the policy increase and more children opt to attend. Both adult and child earnings increase for the same reason. Overall, these results suggest that were the experimental sample not selected, and thus the gains to human capital and schooling observed in the experimental sample could be generalized to the broader population, the GDP gains from the policy would be 3.9 (6.0-2.1) pp larger.

The light orange bars labeled 'Same fertility rate' report the results of the fertility counterfactual where we eliminate the fertility impact for new attendees. Here the gains to GDP, attendance, adults earnings, and child earnings are very similar to the baseline. This is largely driven by the fact that both the gains to SHS attendance in baseline and the estimated fertility impact of attendance are small, thus the overall impact of fertility is small. To the extent that things do change, the GDP gains fall slightly (1.8 percent vs 2.1 percent baseline) as a result of the lower average ability

Figure 4: Effects of Free SHS under Counterfactual Scenarios

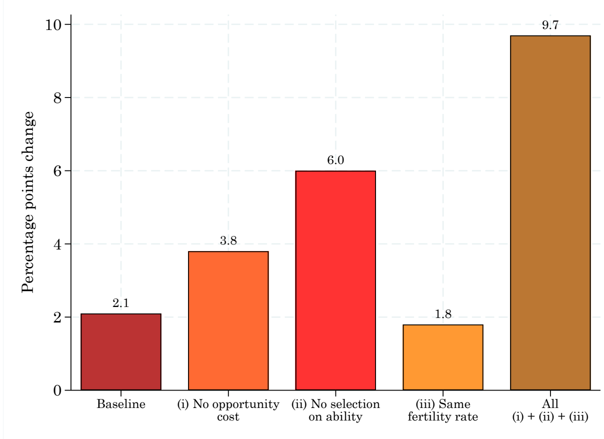

(a) GDP per Capita

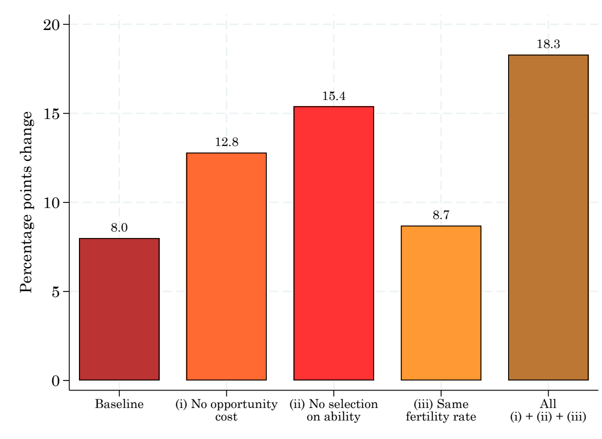

(b) SHS Completion Rate

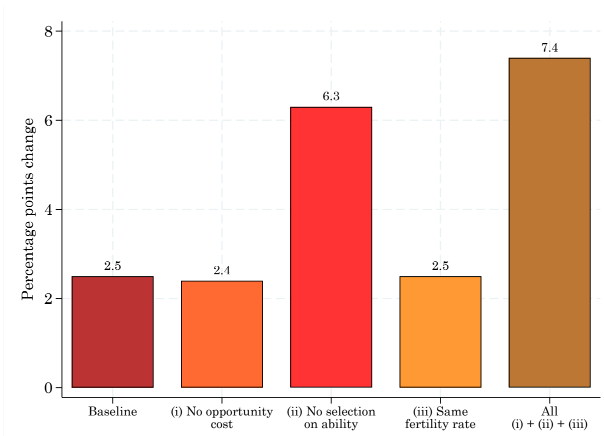

(c) Adult Earnings

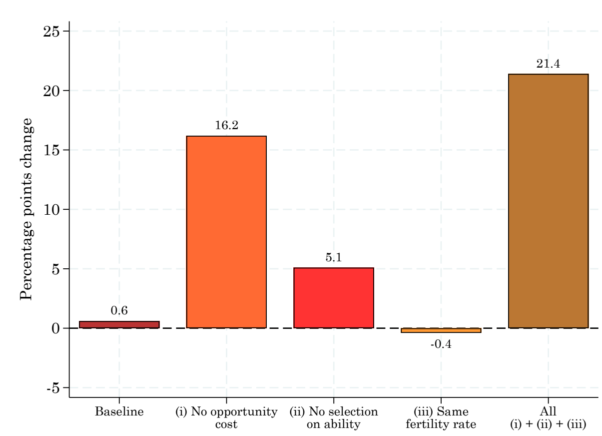

(d) Child Earnings

compared to the baseline. To see why this occurs, Figure 5 displays the change in the long-run ability distribution as a result of the free schooling policy in the baseline model. As shown by the yellow-red positive values for middle-ability parents and children, free schooling leads to an increase in the middle-ability population — this is the group induced to attend by the expansion and, because $\nu_S > \nu_J$, they experience higher net fertility. In response, the population shares of low-ability and high-ability individuals (whose attendance decisions are not impacted by the policy) decline. The net result of these changes is a slight increase in average ability, and thus GDP per capita. Thus the GDP gains are slightly smaller when this channel is shut down.

Figure 5: Changes in Ability Distribution Induced by Free Schooling

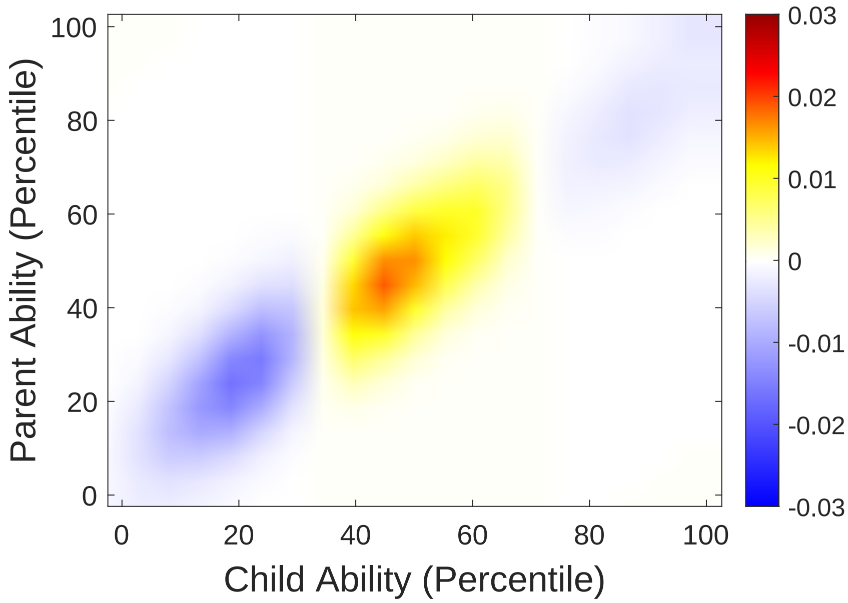

The brown (rightmost) bars of Figure 4 display what happens when we combine all three counterfactuals, shutting down the opportunity cost, selection and differential fertility channels. In this case the GDP gains from free schooling are large, at 9.7 percent, and almost 8 pp higher than in the baseline. Now secondary schooling attendance expands by a robust 18 pp and adult earnings grow by close to 7.4 percent, while child earnings also grow substantially. Though not shown here, it is worth noting that unlike in the baseline analysis, the increase in taxes necessary to pay for the policy (2.7 percent of GDP per capita) are substantially less than the 9.7 percent GDP gains. Thus, although free SHS in the baseline model increases taxes by more than it increases GDP and therefore lowers consumption, the opposite is true here, and consumption grows by about 7 percent. 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 a substantial portion of the gains from higher human capital from free schooling implied by the significant test score and secondary completion increases found by Duflo et al. (2025).

Table 5: GDP Gains and Experimental Outcomes, Alternative Parameterizations

		Treatment Effects		Control Values
	↑ GDP (%)	SHS Attend. (pp)	HC (%)	SHS Attend (%)
Data	-	27	7	48
Estimated model	2.1	27	9	48
+ Schooling quality $\times 2.5$	1.0	16	3	79
+ Schooling cost $\times 2$	3.2	42	19	52
+ Savings wedge $\times 2$	7.7	49	30	48

Note: This table reports the gains in GDP from free secondary schooling under alternative parameter choices, as well as the experimental treatment effects on SHS attendance and human capital measured in these alternative models. Each row changes one additional parameter while maintaining the changes of the rows above it, so that in the fifth row all of schooling quality (η_S), schooling costs (Ψ_S), and the savings wedge (χ) are different from their baseline values.

5.3. The Role of Experimental Moments in Driving Results

While there are straightforward links between certain experimental moments and model parameters (e.g. between the experimental treatment effect on human capital and the parameter governing the return to schooling η_S), it is not immediately clear from the results above how critical the experimental findings are in determining the model's conclusions. The model has many channels and, lacking tractable analytical results, it is difficult to see whether experimental findings could have led to different conclusions. In other words, it is possible that the conclusion of small aggregate gains is baked into the assumptions and design of the model, rather than depending critically on the experimental moments used in estimation.

We perform a quantitative experiment, summarized in Table 5, to investigate this concern. We increase three key parameters of the model that correspond closely to the extent of misallocation and report the gains to GDP from free schooling as well as the experimental moments that *would have been measured* had the experiment taken place in this alternative model. The first two rows display these values in the baseline model and in the data in order to ease comparison.

The third row displays the results when the parameter governing the quality of schooling η_S is increased by a factor of 2.5. The fourth row additionally increases the cost of schooling Ψ_S by a factor of 2 (so that the results in the fourth row correspond to a model where both η_S and Ψ_S have been increased), and the final row doubles the size of the savings wedge.[12] Examining the results, it is clear that the model is able to deliver large GDP gains from the schooling expansion policy. After increasing all three parameters, the GDP gains increase from 2.1 to almost 8 percent. Under these alternative parameters, the return to schooling is large and attendance is low due to high levels of misallocation and, consequently, free schooling leads to large gains. However, this parameterization predicts counterfactually large treatment effects from the experimental evaluation of free schooling, with SHS attendance increasing by 49 pp (vs 27 pp in data) and human capital increasing by 30 percent (vs 8 percent in data).

The upshot of this exercise is that the model is able to deliver, in theory, large GDP gains from free schooling, but that the experimental moments used in model estimation push the model to the conclusion that the gains are, in practice, modest. These results also quantify how different the experimental results would need to be in order for the model to reach a different conclusion. From this perspective, they suggest some minor caution in interpreting our results too broadly — although the treatment effects needed to generate substantial GDP gains are very large, they are not so large that it is impossible to imagine that similar experiments in other contexts and/or locations could find results of this magnitude. If this were to be the case, estimating the model to these contexts would yield very different results from those we find in Ghana.

5.4. Robustness and Discussion

In order to focus on the channels of opportunity cost, selection, and fertility, the model above simplifies away some potentially important features. Here, we discuss some of these features and perform some basic quantitative exercises to assess the robustness of our 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 par-

[12] The factors by which the parameters are increased are chosen so that the control level of SHS attendance after all three parameters have been changed is very similar to that of the baseline model, as can be seen from the fourth column.

ents 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.3 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 (399 percent of GDP per capita vs 73 percent). Similarly, the average valuation for those who do not attend secondary school is smaller than the cost (31 percent of GDP per capita vs 73 percent). 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 that 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) to 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).[13]

Even after incorporating this large externality, implementing the secondary school expansion policy continues to yield results similar to those in the baseline model — the gains from free secondary schooling increase from 2.1 percent of GDP to 3.6 percent. Although the gains are larger, the impact of the policy on GDP remains substantially smaller than the back-of-the-envelope calculation in Section 2 would suggest, and the channels of opportunity cost, selection, and fertility continue to play major roles. The aggregate return to 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.[14]

Rationed Public Sector Jobs: In many developing countries, the ability to obtain

[13]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, suggesting a spillover of 37.5 ($7.5/20$) percent.

[14]Even within this literature on the impacts of human capital, there is a debate over whether it is better understood to impact the level of output (as in the studies cited above) or its growth rate (as in e.g. Hanushek and Woessmann, 2012). Estimates from this latter literature also vary. While we focus on (per capita) levels, it remains possible that substantial spillovers via growth would lead to different conclusions. Indeed, if such effects persist indefinitely, even a small increase in attendance will eventually lead to a large increase in output.

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. (2025) 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,[15] 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 expansion policy are similar to baseline model. Unsurprisingly, the gains to GDP are

[15]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.

smaller than in the baseline model at 0.7 percent (baseline: 2.1 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 gains are only somewhat smaller than the baseline model reflects the fact that the increase in secondary attendance is modest, only about 8 pp and, consequently, the decline in the return to schooling induced by the model extension is small (about 2 percent).

Interest Rate Response: We assume that the economy is small and open, so that interest rates are effectively fixed. While this provides tractability, it means that we may fail to account for shifts in domestic savings behavior. In particular, if paying for schooling is a substantial portion of parents' savings incentive, making schooling free and eliminating this incentive will reduce savings, the capital stock, and thus output. To assess if this is quantitatively significant, we conduct a counterfactual policy exercise where, in addition to making schooling free, we freeze foreign investment at its pre-policy (per capita) level so that any behavioral response to the policy manifests as a change to the capital stock. The results are very similar to our baseline analysis — GDP per capita increases by 1.9 percent (vs 2.1 baseline).

Test Scores and Wages: In reality, a one standard deviation increase in test scores may not lead to exactly one standard deviation increase in earnings (though it is worth noting that it is consistent with the experimental point estimate for earnings). To explore how important this is for our conclusion, we re-estimate the model targeting a weaker relationship. Lacking any evidence from developed countries directly, we choose to match developed country evidence from Murnane, Willett, Duhaldeborde, and Tyler (2000), summarized by Hanushek (2011), where a one SD increase in test scores leads to roughly a one-third SD increase in earnings. With the caveat that the model has some trouble matching this new target and instead delivers a relationship close to one-to-one-half, we find that our results change very little; free schooling increases GDP per capita by 1.7 percent (vs 2.1 percent baseline).

6. Aggregate Effects of Alternative Policies

If making secondary school free to parents does not significantly 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 summarizes the results of various alternative policy counterfactuals. The first row reproduces the key aggregate statistics from the free schooling policy coun-

Table 6: Aggregate Effects of Alternative Policies

	GDP Gain (%)	SHS Increase (pp)
Free Secondary School (Main Analysis)	2.1	8.0
Universal Basic Income	-0.9	-1.6
Tax + Transfer to Bottom 25%	-0.2	-2.5
Allow Borrowing	1.3	8.0
Raise Schooling Quality	9.6	3.6

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).

terfactual from the previous section (i.e. the gains in GDP and the increase in secondary schooling completion). The second and third rows simulate tax-and-transfer schemes. In the presence of substantial misallocation, one might imagine that such policies could substantially improve outcomes by helping to provide cash to poor households, helping them to overcome liquidity constraints and reducing misallocation, without the undesirable effect of distorting parents' and children's' attendance decisions. In both policies we increase the tax rates by the same magnitudes as in the baseline policy but now simply distribute the extra revenues. In the second row, these revenues are used to fund a universal basic income given to all parents. In the third row, these revenues are used to provide transfers only to parent in the bottom 25 percent of the earnings distribution so that the transfers are four times larger. Disappointingly, though not surprising given the results of the previous section, neither of these policies lead to substantial GDP gains (though a planner with a concave social welfare function may still prefer to implement them), and if anything depress GDP slightly as slightly fewer children attend SHS as a result of higher taxes. The fourth row displays the impact of removing the borrowing constraint, allowing parents to borrow to fund their children's education.[16] Here the gains to GDP are modest, reinforcing the intuition from Section 4 that aggregate misallocation is low and low attendance rates are the result of small returns to schooling, rather than

[16]To maintain a coherent concept of equilibrium, we replace the borrowing constraint with a requirement that all individuals die with zero assets so that borrowing does not allow the creation of resources ”out of thin air”.

constraints.

Improving School Quality: 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. This effect size is conservative relative to the average effect 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 substantially larger effects on GDP from improving schooling quality than providing free schooling. GDP increases by 9.6 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 3.6 pp Though this is less than the free schooling policy, the implication is that there is a portion of students who were not attending secondary schooling to begin with because they felt the returns were not high enough to justify the costs (including 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 .

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., 2025). 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).

When we simulate the general equilibrium effects of free secondary school in our model, we find that it raises secondary enrollment substantially but has only a small impact on GDP per capita. In particular, aggregate income increases by much less than back-of-the-envelope calculations based on experimental results would suggest.[17] Three general equilibrium forces serve to drive a wedge between the experimental results and the results produced by the model. 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 changes in fertility for those who attend, altering the distribution of ability.

We conclude that free secondary education policies are mostly redistributive in

[17] 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. (2025) 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.

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 for less fiscal cost and would also work to expand schooling enrollments as it would encourage more enrollment among children for whom the wage gains are marginal at current quality levels. Quantitative exercises show that these conclusions are largely driven by the experimental results used to estimate the model, rather than being an unavoidable result of model mechanisms and/or aggregate data. While this affirms the necessity of leveraging experimental moments to examine the effects of free schooling and other policies, it also warrants some caution in applying our results too broadly as the same experiment could give different results in a different context, particularly in situations or countries where schooling quality is much higher. More broadly speaking, we do not view our results as contradictory to the classic view in development that education is a primary driver of aggregate income. Instead, they further underscore the importance of accounting for education quality in measurement (as in e.g. Schoellman, 2012). Overall, 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.

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.

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 (2025): "The Impact of Secondary School Subsidies on Career Trajectories in a Dual Labor Market: Experimental Evidence from Ghana," .

DUFLO, E., P. DUPAS, E. SPELKE, AND M. P. WALSH (2024): "Intergenerational impacts of secondary education: experimental evidence from Ghana," Tech. rep., National Bureau of Economic Research.

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. (2011): "The economic value of higher teacher quality," *Economics of Education review*, 30, 466–479.

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.

KRUEGER, A. B. AND M. LINDAHL (2001): "Education and Growth: Why and for Whom?" *Journal of Economic Literature*, 39, 1101–36.

KRUEGER, D., A. LUDWIG, AND I. POPOVA (2025): "Shaping Inequality and Intergenerational Persistence of Poverty: Free College or Better Schools?" *Journal of Monetary Economics*, 150.

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

LAGAKOS, D., B. MOLL, T. PORZIO, N. QIAN, AND T. SCHOELLMAN (2018): "Life cycle wage growth across countries," *Journal of Political Economy*, 126, 797–849.

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.

MURNANE, R. J., J. B. WILLETT, Y. DUHALDEBORDE, AND J. H. TYLER (2000): "How important are the cognitive skills of teenagers in predicting subsequent earnings?" *Journal of Policy Analysis and Management*, 19, 547–568.

PORZIO, T., F. ROSSI, AND G. SANTANGELO (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112, 2774–2814.

PRITCHETT, L. (2001): "Where has all the education gone?" *The world bank economic review*, 15, 367–391.

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

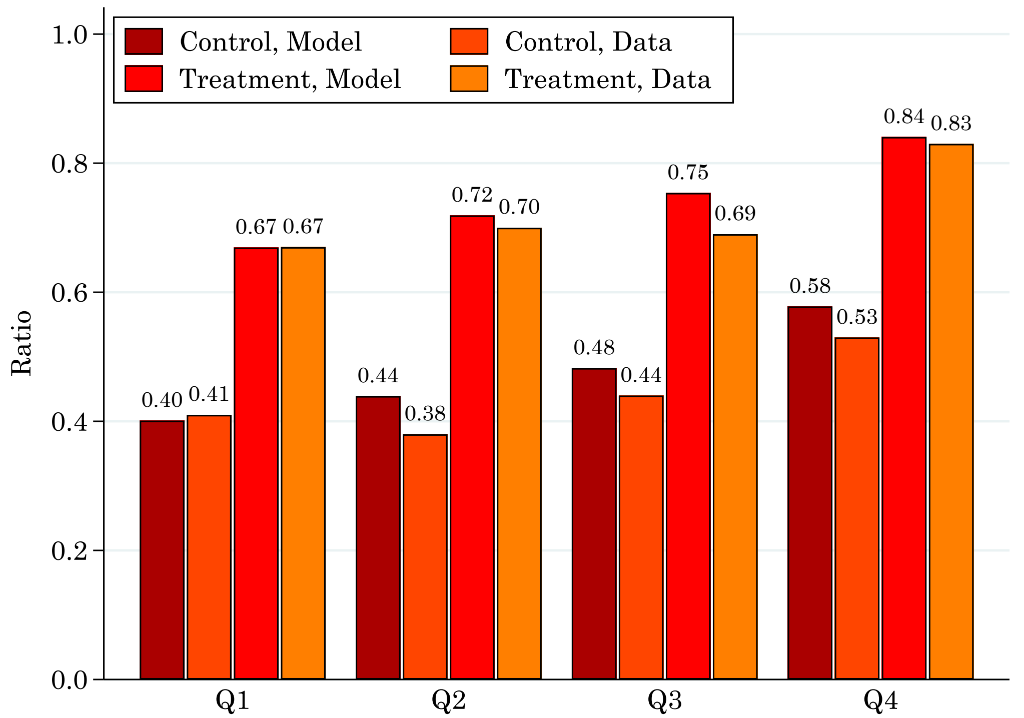

Figure A.2: Total SHS Enrollment in Ghana

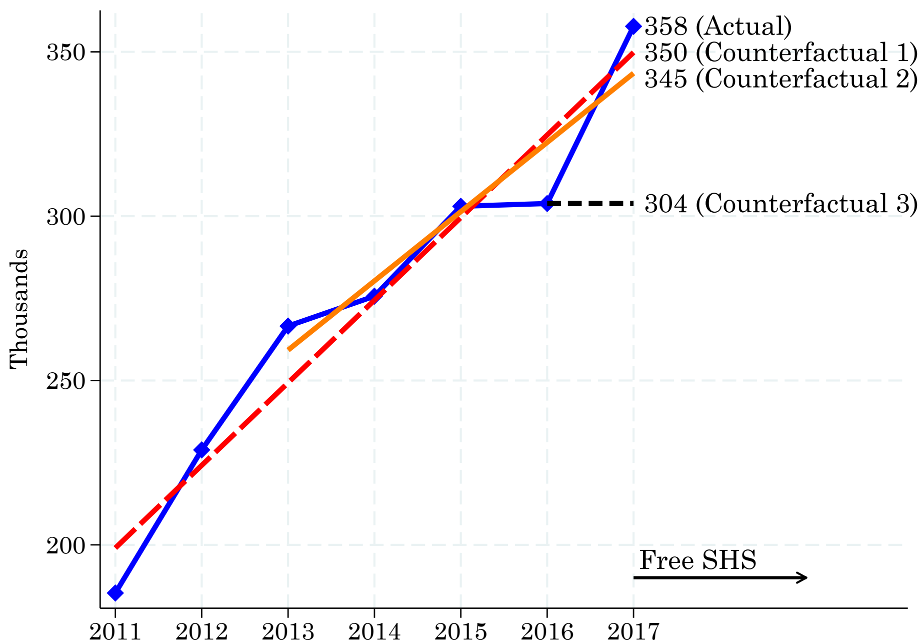

Note: This figure displays the plots corresponding to the various SHS enrollment counterfactuals described in Table A.7. Counterfactual 1 refers to the best-fit line from 2011-2017, Counterfactual 2 refers to the best-fit line from 2013-2017, and Counterfactual 3 refers to the assumption that enrollment levels in 2017 would have been the same as those in 2016 regardless of policy.

Figure A.3: Compensating Variation of Secondary School to Children

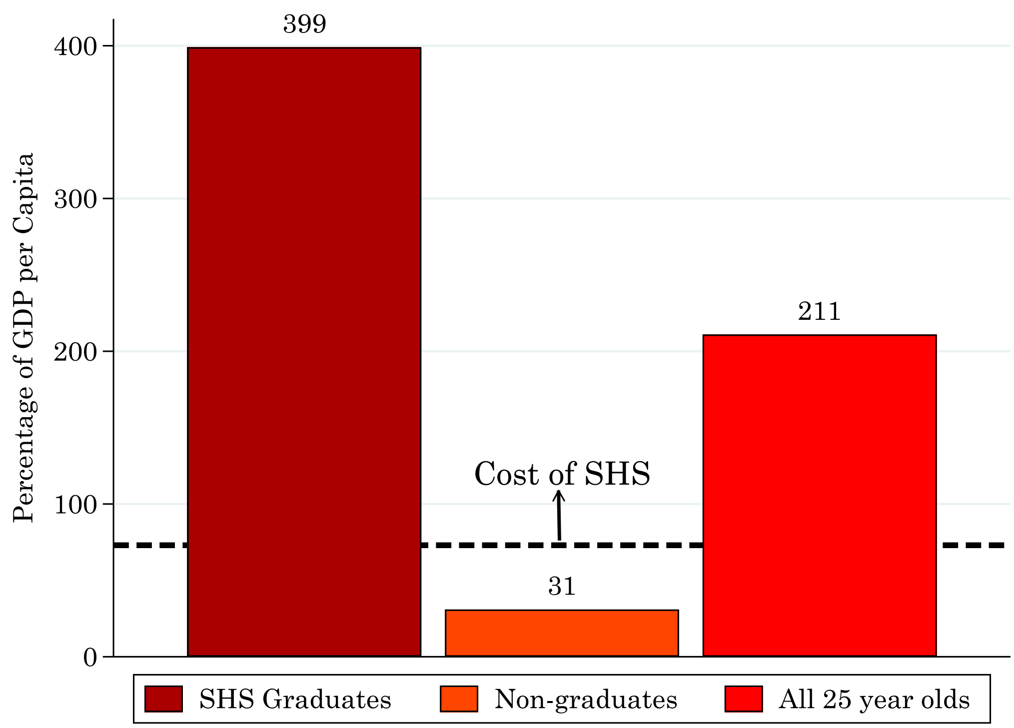

Table A.1: Free Secondary Schooling Policies in Developing Countries

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: Statistics from Figure 2 Economies

	Low Misalloc.	High Misalloc.
Aggregate SHS Attendance	38%	38%
SHS Wage Premium	130%	130%
Treatment Effect on Attend.	+11pp	+50pp
Treatment Effect on Wages	+2.4%	+28%
GDP gains from Free SHS	+0.3%	+12%

Note: This table displays various statistics calculated from the examples economies used to produce Figure 2 and the surrounding discussion.

Table A.3: Control Earnings from Duflo et al. (2025)

Year	Age	Earnings (DDK)	Ghanaian CPI (WB)	Real Earnings
2019	28-29	1458	278.5	524
2023	32-33	3193	610.0	523
Ratio		2.19	2.19	1.0

Note: This table displays nominal earnings in the control group from Duflo et al. (2025) as a function of age, as well as Ghanaian CPI and real earnings.

Table A.4: Employment Rates from 2016/2017 GLSS

	Employed (A)	In School/Training (B)	Neither (A) nor (B) (C)	Employment Rate (A/(A+C))
15-20 years	20.8	63.4	15.8	56.8
20-60 years	79.4	7.1	13.5	85.5
Ratio				0.66

Note: This table displays employment rates by age as computed in the 2016/2017 waves of the GLSS.

Table A.5: 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.

Table A.6: Duflo et al. (2025) Sample vs Nationally Representative Sample

	DDK Control Group	GLSS Rep. Sample
Total Years of Education	11.4	7.6
% Completed SHS	45%	21%
% Self Employed	25%	36%
% w/ Wage Job	35%	26%

Note: This table compares education levels, SHS completion rates, self employment rates, and rates of wage employment between the control group in (Duflo et al., 2025) and a comparable representative sample constructed from the Ghanaian Living Standards Survey. The experimental control group was surveyed in 2019 at an average age of around 28. We use the 2016/2017 waves of the GLSS and limit our sample to 25-30 year olds in order to construct a comparable group.

Table A.7: Actual and Counterfactual SHS Enrollment Increases, 2016 to 2017

	Number	Percent
	Increases in Enrollment	
Actual SHS Enrollment Increase	54,000	17.7
Counterfactual 1 (Trend Since 2011)	46,000	15.1
Counterfactual 2 (Trend Since 2013)	41,000	13.5
Counterfactual 3 (No Change)	0	0
	Increase due to 'Free SHS' (%)	
Counterfactual 1	8,000	2.6
Counterfactual 2	13,000	4.3
Counterfactual 3	54,000	17.7

Note: This table displays the numeric values corresponding to the various SHS enrollment counterfactuals displayed in Figure A.2.

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, the Jacobian matrix is presented in Appendix Table C.1. Here we summarize what we see as the main lessons from this matrix.

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	0.5	0.3	0.0	0.0	0.0	0.1	0.0	0.0	0.0	.00
TE on fertility	-12	13	-0.7	-1.3	0.8	-0.3	0.7	-0.1	0.1	.00
Intergenerational school corr.	0.5	-0.1	0.1	1.6	-0.8	-0.5	0.6	-0.3	0.0	-.01
Var(permanent income)	-0.6	0.6	1.4	4.3	1.6	0.0	0.0	-0.2	0.0	-.01
TE on human capital	0.1	0.0	-0.1	-0.1	0.1	-0.1	0.2	0.0	0.0	.00
TE on SHS completion	0.1	0.2	-0.7	-1.3	0.8	-0.3	0.7	-0.1	0.1	.00
Aggregate SHS attendance	-1.0	0.5	0.5	1.1	1.2	0.6	-0.6	-0.3	0.0	-.01
SHS in top quartile	0.0	0.0	0.2	0.4	-0.1	0.0	0.0	-0.2	0.0	-.01
SHS in bot quartile	0.0	0.2	-0.3	-0.4	0.1	0.0	0.0	0.3	0.0	.00
TE on SHS, Q4-Q1 difference	5.1	-2.3	2.6	6.8	-8.8	-.6	-1.6	8.9	-0.7	.01

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.

D. Welfare Analysis

In this section, we briefly discuss and quantify the impact of the free secondary school on welfare. 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 higher fertility. Thus implementing the policy "creates" some individuals who would have otherwise never been born.

Rather than taking a stance on how to aggregate welfare across agents who may or may not exist depending on 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. To give a sense of the longer run impacts of the policy, we also report the welfare gains for the grandchildren of these parents, restricting our analysis to the set of grandchildren who are always born regardless of policy.

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

Table D.1: Welfare Change of Select Groups under Free Schooling

	Overall	Bottom 25 Pct	Top 25 Pct
Parents	1.7	3.4	-2.5
Children	4.5	18.5	-3.9
Grandchildren (always born)	3.9	-	-

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

top and bottom 25 percent of the (pre-policy) income distribution.

The first two rows of Table D.1 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 1.7 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 3.4 percent of consumption while the richest lose 2.5 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 larger and positive at 4.5 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 3.9 percent, slightly more than the parents, due to higher taxes. The (pre-policy) poorest 25 percent of children, however, make substantial gains of 18.5 percent. While a small portion of these gains occur broadly within this group due to higher unskilled wages, the majority are accrued by the small number of misallocated children who make substantial income gains. The third row of Table D.1 list welfare gains for the grandchildren of the parents described in the first row, limited to those grandchildren who are always born regardless of whether the free schooling policy is implemented or not. They gain 3.9 percent.

Improving School Quality: Table D.2 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 D. For parents, the gains in welfare are similar to those under the free schooling policy (1.4 percent vs 1.7 per-

Table D.2: Welfare Change of Select Groups under Improved Schooling Quality

	Overall	Bottom 25 Pct	Top 25 Pct
Parents	1.4	1.1	2.3
Children	7.5	3.9	14.1
Grandchildren (always born)	6.8	-	-

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

cent), but children and grandchildren benefit substantially more under improved schooling quality. Children's welfare increases by 7.5 percent (vs 4.5 percent baseline) and grandchildren's increases by 6.8 percent (vs 3.9 percent).

The distributional implications of improved quality, however, are substantially different from free schooling. While the poorest children do benefit by 3.9 percent of consumption, this is substantially lower than the 18.5 percent gains from free schooling. Additionally, the largest gains accrue to the richest children whose welfare increased by 14.1 percent compared to -3.9 percent from free schooling. The reason for this substantial switch is twofold. First, children born to richer parents are more likely to attend school, and thus are more like to reap the benefits of improved schooling quality. 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.