GRAPE Working Paper #116

Optimal Taxation of Human Capital with Parental Altruism and Asymmetric Information

Sylwia Radomska and Marek Kapicka

FAME | GRAPE, 2026

Foundation of Admirers and Mavens of Economics
Group for Research in Applied Economics

Optimal Taxation of Human Capital with Parental Altruism and Asymmetric Information

Sylwia Radomska
INE PAN and FAME|GRAPE

Marek Kapicka
CERGE-EI

Abstract

This paper studies optimal education finance in a dynastic Mirrlees economy in which parents derive direct utility from their children's human capital alongside standard dynastic discounting. Education-specific parental altruism adds a non-productive utility return to investment: it raises parental utility independently of the output it generates. We show that this second channel alters the constrained-efficient human-capital wedge: sufficiently strong altruism reverses the wedge from negative to positive, the optimal education subsidy is decreasing in altruism, and stronger altruism shifts intergenerational transfers away from financial bequests toward education. Calibrated to the U.S. economy, the model implies that optimal education support is non-monotonic in income and decreasing in bequests: low-income dynasties receive support due to borrowing constraints, while middle-income families face the weakest case for intervention. Income-contingent loans raise schooling, output, and welfare, but widen educational dispersion. Income-dependent subsidies reduce educational inequality more directly, at the cost of labor-supply distortions and lower aggregate output.

Keywords:
optimal taxation; human capital, parental altruism, asymmetric information, dynastic Mirrlees model, income-contingent loans, education subsidies, intergenerational transfers, bequests.

JEL Classification:
H21, H52, I22, J24, D82, D64

Corresponding author: :
Sylwia Radomska (radomska@inepan.waw.pl).

Acknowledgments
The financial support of National Science Centre (grant UMO-2021/41/N/HS4/02732) is gratefully acknowledged.

Published by: FAME | GRAPE
ISSN: 2544-2473

1 Introduction

Private investment in children's education is large, concentrated among high-income families, and driven partly by preferences that go beyond financial returns. Evidence from surveys, subjective-expectations data, and experiments documents that parents value their children's educational attainment directly—through aspirations, expectations, status, or intrinsic concern for schooling—and that these motives shape educational investments (Bernard et al., 2019; Attanasio and Kaufmann, 2014a; Giannola, 2024). In the United States, private education expenditures account for a non-negligible share of household resources, with substantial dispersion across the wealth distribution. Whether and how this direct parental valuation of education changes optimal education finance is the question this paper addresses.

We study the question in a dynastic Mirrlees economy with endogenous human capital, building on Koeniger and Prat (2018). Each generation chooses consumption, labor supply, financial bequests, and investment in the next generation's human capital. Ability is persistent across generations and privately observed; labor income and educational attainment are observable. The key departure from the existing literature is that parents derive direct utility from their children's acquired human capital—education-specific altruism—over and above its productive value. This preference is distinct from standard dynastic discounting: whereas standard altruism makes parents care about children's future utility, education-specific altruism makes parents value the educational stock itself.

At first glance, this departure might seem inconsequential for policy. Since the planner aggregates the same preferences as private agents, education-specific altruism enters both the social and private objectives symmetrically, and cancels from the first-best conditions. Under private information, however, the picture changes. The planner must satisfy incentive constraints that link education choices to future informational rents. Education-specific altruism raises the private marginal value of human capital and therefore changes the distortions required to implement the constrained optimum—not because the planner disagrees with families, but because altruism alters the allocation of informational rents under asymmetric information.

The main theoretical result is that education-specific altruism alters the sign and magnitude of the constrained-efficient human-capital wedge. In the Koeniger and Prat (2018) benchmark, the human-capital wedge is negative: dynasties underinvest in education relative to the constrained optimum, and the planner implicitly subsidizes it. Adding education-specific altruism raises the private return to education. Once altruism is sufficiently strong, the wedge can reverse sign: altruistic dynasties—particularly wealthy ones—may invest more than the constrained-efficient allocation requires, weakening or reversing the case for marginal education subsidies. The optimal subsidy to education is therefore decreasing in the strength of altruism. A second consequence of altruism is compositional: because education and bequests are both channels of intergenerational transfer but differ in their risk properties, stronger altruism shifts intergenerational transfers away from bequests and toward education. This reallocation feeds back into the incentive structure of the planner's problem.

We characterize these effects analytically through a decomposition of the human-capital wedge into a transfer channel and an incentive channel. Education-specific altruism operates primarily through the transfer channel: by making human capital a more attractive intergenerational transfer vehicle than bequests, it amplifies this component of the wedge and can overturn its sign. The incentive channel—which captures how human capital affects future informational rents through ability persistence—is comparatively stable across calibrations.

Calibrated to the U.S. economy to match schooling attainment, intergenerational earnings persistence, education costs, and moments on bequests and education expenditures, the model delivers sharp quantitative predictions. Optimal education support is non-monotone in income and decreasing in bequests. Low-income dynasties benefit from public support because borrowing constraints bind; middle-income families face the weakest case for intervention; and high-income dynasties may receive less support when altruism-driven overinvestment is concentrated at the top. We evaluate two implementable instruments: income-contingent loans (ICLs) and direct education subsidies. ICLs raise average schooling, output, and welfare by relaxing borrowing constraints and insuring income risk, but can widen educational dispersion because high-ability, high-income families respond more strongly. Income-dependent subsidies reduce educational inequality more directly, at the cost of labor-supply distortions and lower aggregate output.

A key asymmetry emerges from this comparison. Across calibrations, the optimal ICL repayment schedule is nearly invariant to the degree of education-specific altruism, whereas the optimal education subsidy changes qualitatively as altruism increases. This asymmetry arises because ICLs primarily address borrowing constraints and income risk, while education subsidies operate directly on the private marginal cost of human capital investment. In economies where parental altruism is strong, the case for education subsidies is therefore substantially weakened, while the case for income-contingent loans remains robust.

The paper contributes to three literatures. Relative to Koeniger and Prat (2018), we show that education-specific altruism changes not only the level but also the sign and targeting of the human-capital wedge, a result that cannot be obtained by rescaling standard dynastic altruism. Relative to the literature on education subsidies and labor taxation, including Bovenberg and Jacobs (2005) and Stantcheva (2017), we show that the standard case for subsidizing education is weakened, and may be reversed, when parents directly value children's human capital. Relative to the literature on income-contingent loans, we show that parental altruism affects subsidies and loans differently: subsidies operate directly on the margin distorted by altruism, while ICLs primarily address borrowing constraints and income risk.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the model. Section 4 characterizes the constrained-efficient wedges . Section 5 describes the calibration and presents the quantitative results. Section 6 studies the policy experiments. Section 7 concludes.

2 Related Literature

This paper is most closely related to Koeniger and Prat (2018), who study optimal redistribution in a dynastic Mirrlees economy with observable human capital, hidden ability, and intergenerational transfers. In their framework, human capital differs from bequests because its returns are risky and because it affects future informational rents. The present paper builds directly on this environment but introduces education-specific parental altruism: parents derive direct utility from their children's acquired human capital, in addition to standard dynastic discounting. This additional preference channel changes the sign, magnitude, and targeting of the human-capital wedge.

A separate empirical and theoretical literature motivates the assumption that parents value children's human capital directly, and not only through its effect on future earnings. Cunha and Heckman (2007) show that parental investments in children's cognitive and non-cognitive skills exhibit dynamic complementarities and play a central role in human-capital formation. Guryan et al. (2008) document that more educated parents devote significantly more time to child care even after accounting for wages and opportunity costs, while Ramey and Ramey (2009) show that time investments in children by college-educated parents increased sharply over recent decades, consistent with rising concern about educational attainment and college admissions. Evidence based on aspirations and subjective expectations also points to motives beyond narrow financial returns: Bernard et al. (2019) provide experimental evidence that changing parental aspirations affects investments in children's education, and Attanasio and Kaufmann (2014b) show that mothers' and youths' expectations about schooling returns and risks are central determinants of schooling choices. Evidence on parental education further suggests that parental effects on children's schooling are not explained only by income, fertility, or neighborhood characteristics (Chevalier, 2004). More recent survey-experimental evidence shows that parental investments reflect preferences over children's human-capital outcomes, not only resource constraints (Giannola, 2024). On the financial side, Hotz et al. (2023) document that parental wealth and income are strong predictors of college attendance, parental financing decisions, college quality, and graduation outcomes. Finally, Doepke and Zilibotti (2017) provide a theoretical framework in which parents' objectives combine altruism with paternalistic motives, including preferences over children's behavior and outcomes. Taken together, this literature motivates our modeling assumption that parents derive utility directly from children's acquired human capital, a motive distinct from standard dynastic altruism over future consumption.

A closely related literature studies the interaction between education subsidies and labor income taxation. Bovenberg and Jacobs (2005) show that education subsidies and labor taxes should be jointly designed, while Stantcheva (2017) studies optimal taxation with human capital accumulation in dynamic environments and derives conditions under which education expenditures should be subsidized or made tax deductible. The present paper complements this work by showing that the case for subsidizing education is weakened, and may be reversed, when parents directly value children's human capital.

This paper also speaks to the literature on intergenerational mobility and education finance. Work following Becker and Tomes (1979, 1986), Loury (1981), and Galor and Zeira (1993) emphasizes parental resources, borrowing constraints, and the persistence of inequality. More recent quantitative studies, including Restuccia and Urrutia (2004) and Caucutt and Lochner (2020), study the role of parental resources, early investments, and education policy for intergenerational mobility. The present paper differs by studying education finance through the lens of optimal taxation under private information.

Finally, the paper relates to work on income-contingent loans and education finance. Chapman (2006) studies the design and performance of income-contingent loan systems in practice. On the empirical side, Britton and Gruber (2020) find limited adverse effects of income-contingent repayments on labor supply, which supports the view that ICLs may generate smaller labor-supply distortions than income-dependent education subsidies. The present paper contributes to this literature by comparing income-contingent loans with direct education subsidies in a unified dynastic optimal-tax framework.

3 Model

We build on the dynastic Mirrleesian framework of Koeniger and Prat (2018) and introduce two key extensions. First, parents derive direct utility from their children's acquired human capital. Second, we study two education-finance instruments separately—income-contingent loans and education subsidies—and evaluate how each can decentralize the constrained-efficient allocation.

3.1 Family problem: set-up

The economy is populated by altruistic dynasties. In each period, a dynasty consists of one parent and one child. Each parent gives birth to one offspring, so population size is constant over time and normalized to one in every period. Dynasties have a finite planning horizon.

At date t, a dynasty is characterized by parental financial wealth b_t, parental human capital h_t, and parental ability θ_t. In each period, the dynasty observes the parent's ability θ_t and chooses parental labor supply l_t, family consumption c_t, education expenditures e_t invested in the child, and next-period bequests b_{t+1} left to the child.

Preferences are time-separable across generations. Per-period utility is separable in consumption, labor supply, and the child's next-period human capital:

$$U(c_t, l_t, h_{t+1}) = u(c_t) - v(l_t) + \varphi(h_{t+1}), \tag{1}$$

where $u \in C^2(\mathbb{R}_+)$ and $\varphi \in C^2(\mathbb{R}_+)$ are strictly increasing and strictly concave, and $v \in C^2(\mathbb{R}_+)$ is strictly increasing and strictly convex.

Both financial wealth b_t and human capital h_t are publicly observable.

3.2 Ability and information structure

As in Mirrlees (1971), agents differ in ability $\theta_t \in \Theta = [\underline{\theta}, \bar{\theta}]$, which is private information. Ability affects productivity but is not observed by the social planner.

Ability evolves according to a first-order Markov process:

$$\theta_{t+1} \sim F(\cdot \mid \theta_t),$$

where $F(\cdot \mid \theta_t)$ denotes the conditional distribution of next-period ability, with associated density $f(\theta_{t+1} \mid \theta_t)$.

We assume that $f(\theta_{t+1} \mid \theta_t) \in C^2(\Theta \times \Theta)$ and $f(\theta_{t+1} \mid \theta_t) > 0$ for all $(\theta_t, \theta_{t+1}) \in \Theta^2$, so the conditional distribution has full support. We further assume positive persistence, in the sense that $\mathbb{E}[\theta_{t+1} \mid \theta_t]$ is strictly increasing in θ_t.

This specification captures intergenerational persistence in latent ability and earnings capacity, and implies that parental characteristics are informative about the distribution of children's future ability.

3.3 Family problem

Output is produced according to

$$y_t = y(l_t, h_t, \theta_t),$$

where l_t denotes parental labor supply, h_t parental human capital, and θ_t parental ability.

Human capital accumulation follows a Ben-Porath-type technology. If the parent spends e_t on the child's education, the child's next-period human capital is

$$h_{t+1} = H(e_t, h_t),$$

where H is strictly increasing and concave in education expenditures e_t and weakly increasing in parental human capital h_t. We assume that, for each h_t, the function $H(\cdot, h_t)$ is strictly increasing and therefore invertible. Hence, instead of treating education expenditure e_t as the choice variable, we can equivalently let the dynasty choose the child's human capital h_{t+1} directly and write the required education expenditure as

$$e_t = g(h_{t+1}, h_t).$$

The function g is the inverse cost representation of the human-capital production technology. It gives the resource cost of attaining h_{t+1} when parental human capital is h_t. We use this cost representation throughout the planner's problem and the wedge analysis.

Given state (b_t, h_t, θ_t), the dynasty solves

$$W_t(b_t, h_t, \theta_t) = \max_{\{c_t, l_t, h_{t+1}, b_{t+1}\}} \left\{ U(c_t, l_t, h_{t+1}) + \beta \int_{\Theta} W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1}) dF(\theta_{t+1} \mid \theta_t) \right\},$$

subject to

$$c_t + b_{t+1} + g(h_{t+1}, h_t) = (1+r)b_t + y(l_t, h_t, \theta_t).$$

In the terminal period $t = T$, there is no next generation inside the model horizon. Therefore, dynasties do not choose h_{T+1}, education expenditures are zero, and the terminal-period payoff excludes the term $\phi(h_{T+1})$.

Throughout the paper, h_{t+1} denotes the child's acquired human capital, while $g(h_{t+1}, h_t)$ denotes the education expenditure required to attain it. Thus, whenever the planner chooses h_{t+1}, education spending is implicitly given by $e_t = g(h_{t+1}, h_t)$.

3.4 Mechanism design

Information structure and allocations The government cannot observe agents' ability θ_t nor labor supply l_t. In each period, the planner observes dynasty output y_t, consumption c_t, and parental human capital h_t. Since next-period human capital is publicly observable, educational investments are observable through the resulting realization of h_{t+1}.

By the revelation principle, we restrict attention to direct revelation mechanisms in which dynasties report their ability truthfully in each period.

Let $\theta^t \equiv (\theta_0, \theta_1, \ldots, \theta_t)$ denote the history of realized abilities within a dynasty. At date t, an allocation is a measurable mapping $x_t(\theta^t) \equiv \big(c_t(\theta^t),\, y_t(\theta^t),\, h_{t+1}(\theta^t)\big)$. The feasible set $\mathcal{X}$ consists of all sequences $\mathbf{x} \equiv \{x_t(\theta^t)\}_{t=1}^T$ of such measurable functions.

From this point onward, we work with the reduced-form representation of utility obtained after substituting out labor supply. With a slight abuse of notation, we write

$$U(c, y, h, \theta, h') \equiv u(c) - v(y, h, \theta) + \varphi(h').$$

The expected lifetime utility of a dynasty under allocation $\mathbf{x}$ is therefore

$$\mathcal{U}(\mathbf{x}) \equiv \mathbb{E}_0 \left[\sum_{t=1}^{T} \beta^{t-1} U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big) \right].$$

Incentive compatibility Dynasties may misreport their ability through reporting strategies

$$\mathbf{r} \equiv \{r_t(\theta^t)\}_{t=1}^T.$$

An allocation is incentive compatible if truthful reporting yields weakly higher expected utility than any alternative reporting strategy:

$$\mathcal{U}(\mathbf{x}) \geq \mathcal{U}(\mathbf{x} \circ \mathbf{r}) \qquad \text{for all } \mathbf{r} \in \mathcal{R}, \tag{2}$$

where

$$(\mathbf{x} \circ \mathbf{r})(\theta^t) \equiv x_t\big(r^t(\theta^t)\big)$$

denotes the allocation induced by the reporting strategy $\mathbf{r}$.

The planner's objective Following Koeniger and Prat (2018) and Farhi and Werning (2013), the planner minimizes the expected discounted resource cost of providing an allocation subject to incentive compatibility and promise keeping. Let

$$q \equiv (1+r)^{-1}$$

denote the planner's discount factor. Using the cost representation of education, the expected cost of an allocation is

$$\Pi(\mathbf{x}) \equiv \mathbb{E}_0\left[\sum_{t=1}^{T} q^{t-1}\Big(c_t(\theta^t) + g\big(h_{t+1}(\theta^t), h_t(\theta^{t-1})\big) - y_t(\theta^t)\Big)\right].$$

The planner solves

$$\begin{aligned} \min_{\mathbf{x}\in\mathcal{X}} \quad & \Pi(\mathbf{x}) && (3)\\ \text{s.t.} \quad & \mathcal{U}(\mathbf{x}) \geq V, && \text{(i)}\\ & \mathbf{x} \in \mathcal{X}^{IC}. && \text{(ii)} \end{aligned}$$

Condition (i) is the promise-keeping constraint: the dynasty's expected utility under truthful reporting must be at least the promised level V. Condition (ii) restricts attention to incentive-compatible allocations.

A recursive first-order approach A direct solution of (3) is intractable because incentive compatibility must hold against all possible reporting strategies. Following Fernandes and Phelan (2000), Kapička (2013), Farhi and Werning (2013), and Koeniger and Prat (2018), we therefore use a recursive first-order approach. The key advantage is that it replaces the global incentive constraints with local conditions and a recursive envelope condition, thereby allowing the planner's problem to be written recursively without explicitly specifying all off-equilibrium continuation utilities.

Let $\omega_t(\theta^t)$ denote continuation utility after history θ^t. It satisfies

$$\omega_t(\theta^t) \equiv U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big) + \beta \int_\Theta \omega_{t+1}(\theta^t, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t), \qquad (4)$$

with terminal condition

$$\omega_{T+1}(\theta^{T+1}) = 0.$$

At the beginning of period t, dynasties compare the continuation utility under truthful reporting with the utility obtained under alternative reporting strategies. For any strategy $\mathbf{r} \in \mathcal{R}$, define

$$\omega_t^{\mathbf{r}}(\theta^t) \equiv U\big(c_t(r^t(\theta^t)), y_t(r^t(\theta^t)), h_t(r^{t-1}(\theta^{t-1})), \theta_t, h_{t+1}(r^t(\theta^t))\big) + \beta \int_\Theta \omega_{t+1}^{\mathbf{r}}(\theta^t, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t). \qquad (5)$$

Ex-post incentive compatibility requires

$$\omega_t(\theta^t) \geq \omega_t^{\mathbf{r}}(\theta^t) \qquad \text{for all } \theta^t \text{ and all } \mathbf{r} \in \mathcal{R}. \qquad (6)$$

Let $\mathcal{X}^{IC}$ denote the set of allocations satisfying (6). By Theorem 2.1 in Fernandes and Phelan (2000), ex-ante and ex-post incentive compatibility are equivalent up to sets of measure zero in this recursive environment.

Validity of the first-order approach We further restrict attention to one-shot deviations. For recursive dynamic screening problems of this type, Lemma 1 in Kapička (2013) implies that it is without loss of generality to study one-shot deviations, since any profitable multi-period deviation can be decomposed into a sequence of one-shot deviations. Under the regularity and monotonicity conditions stated below, the local incentive constraints together with the recursive envelope condition are sufficient for global incentive compatibility. The formal argument is provided in Appendix F.

Assumption 1 (Regularity, single-crossing, and monotonicity)**.**

1. *Regularity. The functions u, v, φ, and g are twice continuously differentiable. The consumption utility u and the altruism term φ are strictly increasing and strictly concave, the disutility term v is strictly increasing and strictly convex in output y, for given (h, θ), and the education cost function g is strictly increasing in both arguments. The conditional density $f(\theta_{t+1} \mid \theta_t)$ is continuously differentiable and strictly positive on $\Theta \times \Theta$. Continuation utilities $\omega_t(\theta^t)$ are absolutely continuous in θ_t.*

2. *Single-crossing. The reduced-form per-period utility*

$$U(c, y, h, \theta, h') \equiv u(c) - v\big(y, h, \theta\big) + \varphi(h')$$

satisfies increasing differences in (y, θ):

$$\frac{\partial^2 U(c, y, h, \theta, h')}{\partial y\, \partial \theta} = -\frac{\partial^2 v\big(y, h, \theta\big)}{\partial y\, \partial \theta} \geq 0.$$

Equivalently, higher ability weakly reduces the marginal disutility of producing a given output level. Under the production technology

$$y = \theta^{\xi} h^{1-\xi} l, \qquad \xi \in (0, 1),$$

this condition is satisfied. The altruism term $\varphi(h')$ does not affect the single-crossing property, since it depends on future human capital but not on the current type θ.

3. *Monotonicity. The constrained-efficient allocation is weakly monotone in the current type. In particular, for every history θ^{t-1}, both output $y_t(\theta^t)$ and next-period human capital $h_{t+1}(\theta^t)$ are weakly increasing in θ_t. Monotonicity of output follows from standard monotone comparative statics under the single-crossing condition. Monotonicity of next-period human capital is less immediate because the altruism term $\varphi(h_{t+1})$ introduces a non-standard dependence on continuation utilities. We therefore verify this property numerically in the quantitative implementation across all grid points, periods, and values of ϑ considered in the analysis.*

Lemma 1 (Validity of the first-order approach)**.** *Under Assumption 1, any allocation satisfying feasibility, promise keeping, the local incentive constraints, and the recursive envelope condition (9) is globally incentive compatible. In particular, truthful reporting maximizes continuation utility after every history.*

Proof. See Appendix F. □

Relaxed problem We now impose only local incentive constraints and derive a recursive representation of the planner's problem. For any incentive-compatible allocation, consider the one-shot deviation strategy $\mathbf{r}^{\sigma,t}$ indexed by $(\sigma, t) \in \Theta \times \{1, \dots, T\}$, under which the dynasty reports σ in period t and reports truthfully in all other periods:

$$r_s^{\sigma,t}(\theta^s) = \begin{cases} \theta_s, & s \neq t, \\ \sigma, & s = t. \end{cases}$$

Because ability follows a first-order Markov process, the distribution of the child's ability depends only on the parent's reported current type. Hence next-period continuation utility under a one-shot deviation depends on the report rather than on the true type:

$$\omega_{t+1}^{\mathbf{r}^{\sigma,t}}(\theta^t, \theta_{t+1}) = \omega_{t+1}(\theta^{t-1}, \sigma, \theta_{t+1}).$$

Substituting this identity into (5) yields

$$\begin{aligned} \omega_t^{\mathbf{r}^{\sigma,t}}(\theta^t) = U\big(&c_t(\theta^{t-1}, \sigma), y_t(\theta^{t-1}, \sigma), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^{t-1}, \sigma)\big) \\ &+ \beta \int_\Theta \omega_{t+1}(\theta^{t-1}, \sigma, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t). \end{aligned} \tag{7}$$

If incentive compatibility holds, then for almost every history θ^t,

$$\omega_t(\theta^t) = \max_{\sigma \in \Theta} \omega_t^{\mathbf{r}^{\sigma,t}}(\theta^t).$$

Since problem (7) satisfies the assumptions of Theorem 2 in Milgrom and Segal (2002), differentiating with respect to the true current type yields the recursive envelope condition

$$\frac{\partial \omega_t(\theta^t)}{\partial \theta_t} = \frac{\partial U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big)}{\partial \theta_t} + \beta \int_\Theta \omega_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1} \mid \theta_t)}{\partial \theta_t}\, d\theta_{t+1}. \tag{8}$$

Equation (8) describes how continuation utility varies with the true type in an incentive-compatible allocation. The first term captures the direct effect of ability on current utility at the assigned allocation. The second term captures the effect of current reporting on the planner's beliefs about future ability and therefore on future continuation utility. With persistent types, truthful reporting today affects the conditional distribution of the child's ability and thus the future informational rents associated with the dynasty.

Recursive formulation of the relaxed problem With persistent types, promised utility alone is not a sufficient state variable. Following Koeniger and Prat (2018), define the future marginal rent

$$\Phi_{t+1}(\theta^t) \equiv \int_\Theta \omega_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1} \mid \theta_t)}{\partial \theta_t}\, d\theta_{t+1}.$$

Using this definition, (8) becomes

$$\frac{\partial \omega_t(\theta^t)}{\partial \theta_t} = \frac{\partial U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big)}{\partial \theta_t} + \beta \Phi_{t+1}(\theta^t). \tag{9}$$

Let

$$\Gamma(V_t, \Phi_t, h_t, \theta_{t-1}, t)$$

denote the planner's minimum expected continuation cost at date t, conditional on promised utility V_t, promised marginal rent Φ_t, current human capital h_t, and the previous generation's type θ_{t-1}. For $t \geq 2$, the relaxed planner's problem can be written as

$$\Gamma(V_t, \Phi_t, h_t, \theta_{t-1}, t) = \min_{\{c,y,h',V',\Phi'\}} \{\mathit{\Pi}\}, \tag{10}$$

where

$$\mathit{\Pi} \equiv \int_\Theta \Big[c_t(\theta^t) + g\big(h_{t+1}(\theta^t), h_t\big) - y_t(\theta^t) \\ + q\,\Gamma\big(V_{t+1}(\theta^t), \Phi_{t+1}(\theta^t), h_{t+1}(\theta^t), \theta_t, t+1\big)\Big] f(\theta_t \mid \theta_{t-1})\, d\theta_t,$$

subject to

$$\omega_t(\theta^t) \equiv U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big) + \beta V_{t+1}(\theta^t), \tag{11}$$

$$\frac{\partial \omega_t(\theta^t)}{\partial \theta_t} = \frac{\partial U\big(c_t(\theta^t), y_t(\theta^t), h_t(\theta^{t-1}), \theta_t, h_{t+1}(\theta^t)\big)}{\partial \theta_t} + \beta \Phi_{t+1}(\theta^t), \tag{12}$$

$$V_t(\theta^{t-1}) \equiv \int_\Theta \omega_t(\theta^t)\, f(\theta_t \mid \theta_{t-1})\, d\theta_t, \tag{13}$$

$$\Phi_t(\theta^{t-1}) \equiv \int_\Theta \omega_t(\theta^t) \frac{\partial f(\theta_t \mid \theta_{t-1})}{\partial \theta_{t-1}}\, d\theta_t. \tag{14}$$

Equations (13) and (14) define, respectively, promised utility and promised marginal rent as functions of the previous history. Equation (11) links current continuation utility to current allocations and future promised utility. Equation (12) is the recursive envelope equation. As is standard in the optimal-control literature, (12) is understood in integral form.

The recursive problem has five state variables:

$$(V_t, \Phi_t, h_t, \theta_{t-1}, t).$$

Promised utility V_t and promised marginal rent Φ_t summarize continuation incentives, current human capital h_t affects both current production and education costs, and the lagged type θ_{t-1} matters because it determines the conditional distribution of the child's ability. By the Markov property, conditioning on θ_{t-1} is sufficient; past histories affect the problem only through the promised objects (V_t, Φ_t).

We use the inverted cost representation of education throughout the planner's problem. That is, we substitute education expenditures with the equivalent cost function $g(h_{t+1}, h_t)$ obtained from the human-capital accumulation technology. This formulation follows Koeniger and Prat (2018) and is the one used in the Hamiltonian analysis presented in Appendix A.

3.5 Optimality conditions

We now characterize the planner's optimality conditions for human capital under private information. Because ability is privately observed, the first-best allocation with full insurance is generally infeasible: perfect insurance would violate incentive compatibility. The constrained-efficient allocation therefore reflects the familiar trade-off between insurance and incentives.

The derivations below rely on the recursive formulation of the relaxed problem and the associated costate calculations, presented in Appendix A.

Remark 1 (Reciprocal Euler equation). *At the constrained-efficient allocation, the planner's first-order conditions imply*

$$\frac{1}{u_c(c_t)} = \frac{q}{\beta}\,\mathbb{E}_t\left[\frac{1}{u_c(c_{t+1})}\right], \tag{15}$$

where $q \equiv (1+r)^{-1}$ and u_c denotes the marginal utility of consumption.

Proof of Remark 1. See Appendix C. □

Proposition 1 (Optimal human capital condition in the relaxed problem). *The first-order condition with respect to next-period human capital in the relaxed planner's problem (10) is*

$$\begin{aligned} q\Bigg(\int_\Theta \left[\frac{\frac{\partial v(y_{t+1}(\theta_{t+1}),h_{t+1},\theta_{t+1})}{\partial h_{t+1}}}{\frac{\partial u(c_{t+1}(\theta_{t+1}))}{\partial c_{t+1}}} + \frac{\partial g(h_{t+2}(\theta_{t+1}),h_{t+1})}{\partial h_{t+1}}\right] dF(\theta_{t+1}\mid\theta_t) \\ - \int_\Theta \mu_{t+1}(\theta_{t+1})\frac{\partial^2 v(y_{t+1}(\theta_{t+1}),h_{t+1},\theta_{t+1})}{\partial\theta_{t+1}\,\partial h_{t+1}}\,d\theta_{t+1}\Bigg) \\ = \frac{\frac{\partial \varphi(h_{t+1}(\theta_t))}{\partial h_{t+1}}}{\frac{\partial u(c_t(\theta_t))}{\partial c_t}} - \frac{\partial g(h_{t+1}(\theta_t),h_t)}{\partial h_{t+1}}. \end{aligned} \tag{16}$$

The costate variable associated with the incentive constraint is

$$\mu_{t+1}(\theta_{t+1}) = \int_{\underline{\theta}}^{\theta_{t+1}}\left[-\frac{1}{\frac{\partial u(c_t(x))}{\partial c_t(x)}} + \lambda + \gamma\frac{\frac{\partial f(x|\theta_{t-1})}{\partial\theta_{t-1}}}{f(x\mid\theta_{t-1})}\right] dF(x\mid\theta_{t-1}), \tag{17}$$

with boundary conditions

$$\lim_{\theta_{t+1}\to\underline{\theta}}\mu_{t+1}(\theta_{t+1}) = 0, \qquad \lim_{\theta_{t+1}\to\bar{\theta}}\mu_{t+1}(\theta_{t+1}) = 0.$$

The multiplier associated with the promise-keeping constraint (13) is

$$\lambda = \mathbb{E}_t\left[\frac{1}{\frac{\partial u(c_t(\theta_t))}{\partial c_t}}\right],$$

while γ is the multiplier associated with the marginal-rent recursion (14).

Proof of Proposition 1. See Appendix C. □

Proposition 1 can be rewritten as

$$\begin{aligned}\frac{\partial g(h_{t+1}(\theta_t), h_t)}{\partial h_{t+1}} =& \frac{\frac{\partial \varphi(h_{t+1}(\theta_t))}{\partial h_{t+1}}}{\frac{\partial u(c_t(\theta_t))}{\partial c_t}} \\ &- q\int_{\Theta}\left(\frac{\frac{\partial v(y_{t+1}(\theta_{t+1}), h_{t+1}, \theta_{t+1})}{\partial h_{t+1}}}{\frac{\partial u(c_{t+1}(\theta_{t+1}))}{\partial c_{t+1}}} + \frac{\partial g(h_{t+2}(\theta_{t+1}), h_{t+1})}{\partial h_{t+1}}\right) dF(\theta_{t+1} \mid \theta_t) \\ &+ q\int_{\Theta} \mu_{t+1}(\theta_{t+1}) \frac{\partial^2 v(y_{t+1}(\theta_{t+1}), h_{t+1}, \theta_{t+1})}{\partial \theta_{t+1}\, \partial h_{t+1}}\, d\theta_{t+1}. \end{aligned} \tag{18}$$

Equation (18) characterizes the planner's trade-off between the marginal cost and the marginal benefit of investing in children's human capital. The right-hand side contains four conceptually distinct components.

First, parental altruism generates a direct utility benefit from higher human capital of the child. This is the new force relative to Koeniger and Prat (2018). Because parents value both current consumption and their child's future human capital, the planner internalizes an additional intradynastic substitution margin between current resources and educational investment.

Second, higher human capital relaxes future production distortions by reducing the disutility of generating a given level of output.

Third, human capital affects the cost of future human-capital accumulation through the dependence of $g(h_{t+2}, h_{t+1})$ on parental human capital.

Fourth, human capital affects incentive compatibility by changing future informational rents, as captured by the term involving μ_{t+1} and the cross-partial derivative of the disutility term with respect to ability and human capital.

4 The wedges

We compare the optimality conditions in the laissez-faire economy to those characterizing the constrained-efficient allocation under private information. The wedges between these conditions capture the implicit marginal taxes or subsidies required to implement the social optimum.

The key economic friction is asymmetric information about ability. While dynasties privately observe their productivity, the planner cannot condition allocations on true ability. As a result, the constrained-efficient allocation distorts behavior relative to the laissez-faire benchmark in order to relax incentive constraints. These distortions are summarized by wedges on savings, labor supply, and human capital investment.

4.1 Laissez-faire problem

In the laissez-faire economy, each dynasty takes prices as given and solves a dynamic problem under private information about its own ability. Conditional on history θ^{t-1}, the dynasty chooses allocations as functions of current ability θ_t:

$$\{c_t(\theta^t),\, y_t(\theta^t),\, h_{t+1}(\theta^t),\, b_{t+1}(\theta^t)\}.$$

The value function satisfies

$$W_t(b_t, h_t, \theta_{t-1}) = \max_{\{c_t, y_t, h_{t+1}, b_{t+1}\}} \int_{\Theta} \Big[U\big(c_t(\theta^t), y_t(\theta^t), h_t, \theta_t, h_{t+1}(\theta^t)\big) + \beta W_{t+1}\big(b_{t+1}(\theta^t), h_{t+1}(\theta^t), \theta_t\big)\Big]\, dF(\theta_t \mid \theta_{t-1}), \tag{19}$$

subject to the budget constraint

$$c_t(\theta^t) + b_{t+1}(\theta^t) + g\big(h_{t+1}(\theta^t), h_t\big) = (1+r)b_t + y_t(\theta^t). \tag{20}$$

There are no taxes or subsidies in the laissez-faire economy.

4.2 Laissez-faire optimality conditions

The laissez-faire allocation is characterized by the following first-order conditions.

Remark 2 (Laissez-faire optimality conditions)**.** *For each history θ^t, the optimal allocation satisfies*

$$\frac{\partial u(c_t(\theta^t))}{\partial c_t} = \beta(1+r)\,\mathbb{E}_t\left[\frac{\partial u\big(c_{t+1}(\theta^{t+1})\big)}{\partial c_{t+1}}\right], \tag{21}$$

$$\frac{\partial u(c_t(\theta^t))}{\partial c_t}\frac{\partial g(h_{t+1}(\theta^t), h_t)}{\partial h_{t+1}} = \frac{\partial \varphi(h_{t+1}(\theta^t))}{\partial h_{t+1}} + \beta\mathbb{E}_t\left[\frac{\partial u\big(c_{t+1}(\theta^{t+1})\big)}{\partial c_{t+1}}\left(\frac{\partial y_{t+1}(\theta^{t+1})}{\partial h_{t+1}} - \frac{\partial g\big(h_{t+2}(\theta^{t+1}), h_{t+1}(\theta^t)\big)}{\partial h_{t+1}}\right)\right], \tag{22}$$

$$\frac{\partial u(c_t(\theta^t))}{\partial c_t}\frac{\partial y_t(\theta^t)}{\partial l_t} = \frac{\partial v(l_t(\theta^t))}{\partial l_t}. \tag{23}$$

Equation (21) is the standard Euler equation for intergenerational transfers. Equation (22) equates the marginal cost of human capital investment to its marginal benefit, which includes both expected future returns and the direct utility from children's human capital due to parental altruism. Equation (23) is the standard intratemporal condition.

Definition of wedges. Wedges measure the gap between private optimality conditions in the laissez-faire equilibrium and the corresponding conditions in the constrained-efficient allocation. Each wedge can be interpreted as the implicit marginal tax or subsidy that decentralizes the planner's allocation.

Formally, for each history θ^t, we define wedges on bequests, labor supply, and human capital investment as the proportional distortions that reconcile the laissez-faire first-order conditions with their constrained-efficient counterparts. These wedges are defined so that the distorted laissez-faire first-order conditions coincide with the planner's optimality conditions.

Definition 1 (Wedges)**.** *For each history θ^t, define:*

$$\tau_b(\theta^t) \equiv 1 - \frac{q}{\beta} \frac{\frac{\partial u\left(c_t(\theta^t)\right)}{\partial c_t}}{\mathbb{E}_t\left[\frac{\partial u(c_{t+1}(\theta^{t+1}))}{\partial c_{t+1}}\right]}, \tag{24}$$

$$\tau_l(\theta^t) \equiv 1 - \frac{\frac{\partial v\left(l_t(\theta^t)\right)}{\partial l_t}}{\frac{\partial u(c_t(\theta^t))}{\partial c_t}\frac{\partial y_t(\theta^t)}{\partial l_t}}, \tag{25}$$

$$\tau_h(\theta^t) \equiv \frac{\beta}{\frac{\partial g(h_{t+1}(\theta^t),h_t)}{\partial h_{t+1}}} \mathbb{E}_t\left[\frac{\frac{\partial u\left(c_{t+1}(\theta^{t+1})\right)}{\partial c_{t+1}}}{\frac{\partial u(c_t(\theta^t))}{\partial c_t}}\left(\frac{\partial y_{t+1}(\theta^{t+1})}{\partial h_{t+1}} - \frac{\partial g\left(h_{t+2}(\theta^{t+1}), h_{t+1}(\theta^t)\right)}{\partial h_{t+1}}\right)\right]$$
$$+ \frac{\frac{\partial \varphi\left(h_{t+1}(\theta^t)\right)}{\partial h_{t+1}}}{\frac{\partial u(c_t(\theta^t))}{\partial c_t}\cdot\frac{\partial g(h_{t+1}(\theta^t),h_t)}{\partial h_{t+1}}} - 1. \tag{26}$$

A positive wedge corresponds to an implicit marginal tax, while a negative wedge corresponds to a subsidy. These wedges summarize how the planner distorts behavior in order to balance insurance against incentives.

The bequest wedge captures intertemporal distortions arising from private information and is positive due to Jensen's inequality, reflecting the cost of providing insurance over time.

The labor wedge captures distortions in production decisions. By reducing output relative to the first-best, the planner limits informational rents that would otherwise accrue to high-ability dynasties.

The human capital wedge is more complex and reflects both intertemporal and incentive considerations. In particular, investment in human capital affects not only future productivity but also the evolution of private information and the tightness of incentive constraints.

Constrained-efficient wedges. The wedges implied by the planner's optimality conditions take standard forms for bequests and labor supply, while the human capital wedge extends these results to incorporate parental altruism.

Proposition 2 (Constrained-efficient bequest wedge)**.** *In the constrained-efficient allocation,*

$$\tau_b^*(\theta^t) = 1 - \frac{1}{\mathbb{E}_t\left[\frac{1}{\frac{\partial u(c_{t+1}(\theta^{t+1}))}{\partial c_{t+1}}}\right]\mathbb{E}_t\left[\frac{\partial u(c_{t+1}(\theta^{t+1}))}{\partial c_{t+1}}\right]}, \tag{27}$$

and $\tau_b^*(\theta^t) > 0$.

Proof. See Appendix C. □

The bequest wedge is strictly positive by Jensen's inequality, since the function $1/u_c(\cdot)$ is strictly convex under standard assumptions, implying that the product of expectations exceeds one. Intuitively, reducing intergenerational transfers lowers the planner's future incentive costs, since higher inherited resources weaken the sensitivity of continuation allocations to descendants' own realizations and choices.

Proposition 3 (Constrained-efficient labor wedge)**.** *In the constrained-efficient allocation,*

$$\tau_l^*(\theta^t) = -\frac{\mu_t(\theta^t)}{f(\theta_t \mid \theta_{t-1})}\frac{\partial^2 v\left(y_t(\theta^t), h_t, \theta_t\right)}{\partial \theta_t\, \partial y_t}. \tag{28}$$

Proof. See Appendix C. □

The sign of the labor wedge is determined by the cross-partial $\partial^2 v/\partial\theta\,\partial y$. Under Assumption 1, higher ability reduces the marginal disutility of producing output, so $\partial^2 v/\partial\theta\,\partial y < 0$. Whenever the costate $\mu_t(\theta^t)$ is positive — which occurs when the incentive constraint binds — the labor wedge is strictly positive.

The economic interpretation is that dynasties do not internalize the effect of higher output on the tightness of incentive constraints: a more productive dynasty increases informational rents for lower types, making redistribution more costly. The planner therefore distorts production relative to the laissez-faire in order to reduce these rents. The magnitude of the distortion is governed by the ratio $\mu_t(\theta^t)/f(\theta_t \mid \theta_{t-1})$, which captures the local severity of the incentive problem.

Proposition 4 (Constrained-efficient human capital wedge)**.** *Suppose output is produced according to $y = A(\theta, h)\,l$, where $A(\theta, h)$ is twice continuously differentiable, and let $\boldsymbol{v} : \mathbb{R}_+ \to \mathbb{R}$ be twice continuously differentiable and strictly increasing. Then, at the constrained-efficient allocation, the wedge on human capital investment admits the decomposition*

$$\tau_h^* = \Delta_b + \Delta_i, \tag{29}$$

where

$$\begin{aligned}\Delta_b \equiv\ & \frac{q}{\frac{\partial g(h',h)}{\partial h'}} \frac{\tau_b^*}{1-\tau_b^*} \mathbb{E}\left[\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right] \\ & + \frac{\beta}{\frac{\partial g(h',h)}{\partial h'}\,\frac{\partial u(c)}{\partial c}} \operatorname{Cov}\left(\frac{\partial u(c')}{\partial c'}, \frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right),\end{aligned} \tag{30}$$

$$\Delta_i \equiv -\frac{q}{\frac{\partial g(h',h)}{\partial h'}} \mathbb{E}\left[\Xi(\theta',h')\,l'(\theta')\,\frac{d\boldsymbol{v}(l'(\theta'))}{dl'}\,\frac{\frac{\partial A(\theta',h')}{\partial\theta'}\frac{\partial A(\theta',h')}{\partial h'}}{A(\theta',h')^2}\,\frac{\mu'(\theta')}{f(\theta'|\theta)}\right], \tag{31}$$

and

$$q\,\frac{\tau_b^*}{1-\tau_b^*} = \mathbb{E}\left[\beta\,\frac{\frac{\partial u(c')}{\partial c'}}{\frac{\partial u(c)}{\partial c}} - q\right], \tag{32}$$

$$\Xi(\theta,h) \equiv 1 - \frac{A_{\theta h}(\theta,h)\,A(\theta,h)}{A_\theta(\theta,h)\,A_h(\theta,h)}. \tag{33}$$

All expectations are taken with respect to $F(\theta'|\theta)$.

Proof. See Appendix C. □

Remark 3 (Economic interpretation)**.** ***The transfer channel*** **Δ_b.** *The term Δ_b captures the interaction between human capital investment and intergenerational transfers. Both bequests and education allow parents to transfer resources to their children, but they are imperfect substitutes because the return to human capital depends on the child's realized ability. The first term in Δ_b scales the expected net return to education by the normalized bequest wedge $\tau_b^*/(1-\tau_b^*)$. The second term is a risk adjustment, reflecting the covariance between the marginal utility of future consumption and the net return to human capital. Under standard conditions, this covariance is negative, so human capital is a poorer hedge against consumption risk than bequests. The sign of Δ_b is therefore generally ambiguous; however, when the expected-return term dominates—as in our calibration—Δ_b is positive.*

The incentive channel Δ_i***.*** *The term* Δ_i *captures the incentive effect of human capital through the persistence of private information across generations. Its sign depends on* $\Xi(\theta', h')$*. When* $\Xi(\theta', h') > 0$*, higher human capital weakens the productivity advantage of high-ability children, thereby relaxing future incentive constraints and pushing* $\Delta_i < 0$*. When* $\Xi(\theta', h') < 0$*, higher human capital strengthens the productivity advantage of high-ability children, tightening future incentive constraints and pushing* $\Delta_i > 0$*.*

The role of parental altruism. *The parameter* ϑ*, governing direct utility from children's human capital, does not appear explicitly in* Δ_b *or* Δ_i*, because* $\varphi(h')$ *enters symmetrically in the laissez-faire condition and in the planner's optimality condition and therefore cancels in the wedge decomposition.*

Nevertheless, parental altruism affects τ_h^* *indirectly through its impact on the constrained-efficient allocation. First, it shifts the consumption path and therefore the bequest wedge* τ_b^**, affecting the transfer component* Δ_b*. Second, it alters the costate variables* $\mu'(\theta')$*, thereby modifying the incentive component* Δ_i*. Third, it changes the joint distribution of* (c', y', h')*, affecting the expectations and covariance terms in* Δ_b*.*

As a result, higher altruism raises the private marginal value of human capital and reduces the gap between private and social incentives. When ϑ *is sufficiently large, private incentives may exceed the social optimum, and the constrained-efficient human capital wedge* τ_h^* *can become positive, rationalizing an implicit tax rather than a subsidy on education investment.*

Lemma 2 (Cross-partial identity)**.** *Let*

$$y = A(\theta, h)\, l \qquad \text{and} \qquad v(y, h, \theta) = \boldsymbol{v}\left(\frac{y}{A(\theta, h)}\right).$$

Then

$$\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial h} = -\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial y}\frac{\partial y}{\partial h} + \frac{d\boldsymbol{v}(l)}{dl}\frac{l}{A(\theta,h)^2}\frac{\partial A(\theta,h)}{\partial\theta}\frac{\partial A(\theta,h)}{\partial h}\Xi(\theta,h), \tag{34}$$

where

$$\Xi(\theta,h) = 1 - \frac{A_{\theta h}(\theta,h)\,A(\theta,h)}{A_\theta(\theta,h)\,A_h(\theta,h)}.$$

Proof. See Appendix C. □

The decomposition in (29) reveals two distinct forces. The term Δ_b captures the intertemporal-transfer channel: because human capital is a risky transfer technology, its valuation differs from that of bequests through the covariance between its return and future marginal utility. The term Δ_i captures the implementability channel: human capital affects future informational rents by changing how strongly productivity responds to ability.

Relative to Koeniger and Prat (2018), parental altruism does not introduce a separate additive term in the decomposition under the present wedge definition. Instead, it changes the equilibrium values of both components through its effect on the constrained-efficient allocation. In particular, higher altruism raises the private marginal value of human capital and can therefore overturn the sign of the wedge, as documented quantitatively in Section 5.

Altruism does not eliminate distortions. Dynasties internalize their descendants' utility, but not the effect of their choices on the implementability constraints faced by the planner. This implies that optimal education policy cannot be studied in isolation: distortions in intergenerational transfers and labor supply interact with human capital investment in determining the net return to education.

More broadly, the model uncovers a novel trade-off between production efficiency and incentive provision. While human capital accumulation relaxes production distortions, it may simultaneously tighten incentive constraints by increasing informational rents in environments with persistent private information. The quantitative importance of these channels and their interaction with parental altruism is assessed in Section 5.

5 Quantitative Model

Having characterized the constrained-efficient wedges analytically, we now assess their quantitative relevance by calibrating the model to the U.S. economy. The quantitative analysis serves two purposes. First, it disciplines the model's structural parameters. Second, it allows us to evaluate how the degree of parental altruism ϑ shapes the sign and magnitude of the human capital wedge.

To solve the dynastic optimization problem, we employ the endogenous gridpoint method (EGM), originally proposed by Carroll (2006) and extended by Hintermaier and Koeniger (2010). This approach is well suited for stochastic dynamic models with continuous state variables and controls, and allows for an efficient and accurate characterization of policy functions without resorting to value function iteration.

The extension developed by Hintermaier and Koeniger (2010) is particularly important in our setting, as it accommodates inequality constraints that may bind only occasionally and across multiple endogenous choices. This feature is essential for capturing borrowing limits and corner solutions in education and bequest decisions that arise endogenously in dynastic environments.

The model is calibrated to the U.S. economy. The joint distribution of initial wealth and education is constructed using data from the 2004 Survey of Consumer Finances (SCF). In particular, we approximate the empirical joint distribution by matching average years of schooling and net worth across net-worth quartiles. Ability follows a stationary Markov process consistent with intergenerational persistence estimates in the literature.

We simulate the economy using 200,000 dynasties drawn from the stationary joint distribution of ability, wealth, and human capital. Each dynasty is followed over four generations, corresponding to a 120-year horizon. To abstract from transitional dynamics and initial-condition effects, we focus on the second generation when evaluating model implications and comparing simulated outcomes to their empirical counterparts.

The numerical solution of the household problem relies on the first-order conditions derived in the previous sections. Policy functions for consumption, labor supply, education investment, and bequests are computed using backward induction. Details of the algorithm and implementation are provided in Appendix D.

5.1 Household Optimization

In the decentralized economy each dynasty solves the following dynamic optimization problem:

$$\hat{W}(b_t, h_t, \theta^t) = \max_{\{c_t, l_t, h_{t+1}, b_{t+1}\}} \left\{ U(c_t, l_t, h_{t+1}) + \beta \int_{\Theta} \hat{W}(b_{t+1}, h_{t+1}, \theta_{t+1})\, dF(\theta_{t+1}|\theta_t) \right\}, \tag{35}$$

subject to the budget constraint

$$b_{t+1} + c_t + g(h_{t+1}, h_t) = (1 + r_t)b_t + y(l_t, h_t, \theta_t) - T(b_t, y_t, h_{t+1}, \theta_{t-1}), \tag{36}$$

the borrowing constraint

$$b_{t+1} \geq \max\left\{ -\phi\, g(h_{t+1}, h_t), \underline{b}_{t+1} \right\}, \tag{37}$$

and the feasibility constraint on human capital,

$$h_{t+1} \geq 1. \tag{38}$$

Output is produced according to

$$y(l_t, h_t, \theta_t) = \theta_t^{\xi} h_t^{1-\xi} l_t, \tag{39}$$

where θ_t denotes innate ability and $\xi \in (0, 1)$ governs the elasticity of output with respect to ability.

The tax system is as defined in equation (40), written in period-t variables. Families are allowed to finance a fraction $\phi \in (0, 1)$ of education expenditures through debt, reflecting the availability of student loans that are repaid by the next generation. Borrowing is additionally limited by an exogenously given debt ceiling $\underline{b}_{t+1}$, as specified in equation (37).

The first-order conditions of the decentralized household problem coincide with those of the laissez-faire benchmark derived in Appendix B. Policy instruments affect household behavior exclusively through wedges induced by the tax system, without altering the underlying structure of optimality conditions.

5.2 Calibration and Quantitative Environment

The calibration pins down the key preference, technology, and information parameters that govern intergenerational investment, with particular emphasis on matching moments related to education, mobility, and wealth transmission. In the calibration, we set the fraction of education expenditures that can be financed through borrowing to $\phi = 0.5$ and impose an exogenous debt limit $\underline{b}' = -\$30{,}000$. This specification allows families to finance up to \$30,000 of investment in their children's human capital through student loans repaid by the next generation. The chosen value is consistent with empirical evidence on average student debt among U.S. college graduates.

The model incorporates a stylized version of the U.S. tax and transfer system in order to evaluate the quantitative implications of alternative education financing regimes. The implementation closely follows the approach in Koeniger and Prat (2018). Taxes are levied separately on labor income, bequests, and education expenditures, and are summarized by

$$T(b_t, y_t, h_{t+1}, \theta_{t-1}) = T^b(b_t) + T^y(y_t) + T^h(h_{t+1}, h_t). \tag{40}$$

Unless stated otherwise, we set $T^h(h', h) = 0$, which allows us to interpret $g(h', h)$ as the net cost of accumulating children's human capital. Labor income taxes are approximated using the parametric form proposed by Heathcote et al. (2017) for the U.S. economy,

$$T^y(y_t) = y_t - \delta y_t^{1-t_y}, \tag{41}$$

capturing the progressivity of the income tax schedule.

Preferences. Each generation in the model corresponds to 30 years. The altruistic discount factor in the family problem is set to $\beta = 0.17$ to match a bequest–wealth ratio of 0.008, as documented by De Nardi and Yang (2016). This parameter therefore reflects intergenerational discounting rather than an annual rate.

The social planner's discount factor is set to $q = 0.412$, which corresponds to an annual discount rate of 3 percent compounded over 30 years, i.e. $q = (1.03)^{-30}$.

To distinguish between parents' willingness to transfer resources through bequests and their direct valuation of children's education, we introduce a separate altruism parameter ϑ governing utility from children's human capital. The per-period utility function is

$$U(c_t, l_t, h_{t+1}) = \ln c_t - \frac{l_t^\alpha}{\alpha} + \vartheta \ln h_{t+1}.$$

Following Chetty (2012), we target a Frisch elasticity of labor supply equal to 0.5, which implies $\alpha = 3$. The altruism parameter is calibrated to $\vartheta = 5 \times 10^{-4}$ to match the empirical fact that families spend approximately 2 percent of lifetime income on education.

Human capital accumulation. The cost of investing in children's human capital follows a Ben-Porath-type technology,

$$g(h_{t+1}, h_t) = \kappa \left[\ln h_{t+1}\right]^{\varsigma_1} \left[\ln h_t\right]^{\varsigma_2}, \tag{42}$$

where h_{t+1} denotes the child's human capital and h_t parental human capital. The parameter κ is calibrated to match average years of schooling equal to 12.8, based on Barro and Lee (2013). The curvature parameters ς_1 and ς_2 are chosen to match, respectively, the average net cost of one year of education and the intergenerational correlation in educational attainment. This formulation captures the idea that higher parental education reduces the effective cost of investing in children's skills, reflecting intergenerational transmission of knowledge beyond formal schooling.

Production technology and ability process. Output is produced according to

$$y_t = l_t \theta_t^\xi h_t^{1-\xi}, \tag{43}$$

where θ_t denotes innate ability and $\xi \in (0, 1)$. This specification implies log-linear wages,

$$\ln w(\theta_t, h_t) = (1 - \xi) \ln h_t + \xi \ln \theta_t,$$

which is consistent with Mincerian wage regressions. The term $(1 - \xi) \ln h$ captures returns to schooling, while $\xi \ln \theta$ represents residual wage dispersion after conditioning on education.

Ability θ_t is assumed to follow a log-normal AR(1) process, consistent with the estimates in Heathcote et al. (2017). Although the social planner observes output y_t, it cannot infer labor supply directly due to the presence of privately observed and stochastic ability.

Targeted and empirical moments used in the calibration are reported in Table 2. Average annual education expenditures in the United States amount to approximately \$21,300, with about 35 percent subsidized by the government, implying a net private cost of \$13,845 (Stantcheva, 2015, 2017).

In the quantitative analysis, we focus on dynasties whose initial generation consists of high-school graduates. For normalization, parental productivity in the initial generation is set to $\theta_0 = 1$.

Table 2 reports targeted statistics in the data and in the model under different levels of education-specific altruism. The baseline calibration provides a close fit to key moments on schooling, intergenerational persistence, and education expenditures. The altruism parameter $\vartheta = 0.0005$ implies that the marginal utility weight on human capital is small relative to consumption. A one log-point increase in human capital yields roughly 0.05 percent of the marginal utility of a one log-point increase in consumption. This is consistent with the calibration, where education expenditures represent a small share of total resources.

Table 1: Calibrated Parameter Values

Parameter	**Target**	**Source**
Preferences		
Discount / altruistic factor $\beta = 0.17$	Bequest–wealth ratio 0.0088	De Nardi and Yang (2016)
Discount factor $q = 0.412$	Annualized interest rate 3%	Standard
Disutility of labor $v(l) = \frac{l^\alpha}{\alpha}$, $\alpha = 3$	Frisch elasticity = 1/2	Chetty (2012)
Altruism toward children's education $\vartheta = 0.0005$	Education expenditures equal to 2% of lifetime income	High School Longitudinal Study of 2009 (HSLS:09)
Production Technology		
Output per unit of labor $y/l = \theta^\xi h^{1-\xi}$, $\xi = 0.9$	Returns to education 10%	Card (1999)
Ability Process		
AR(1): $\ln\theta' = \rho\ln\theta + \varepsilon$		
$\ln\theta' \sim \mathcal{N}\left(-\frac{\sigma^2}{2(1-\rho^2)}, \frac{\sigma^2}{1-\rho^2}\right)$		
Persistence $\rho = 0.459$	Intergenerational earnings elasticity $\psi = 0.45$	Chetty et al. (2014)
Variance $\frac{\sigma^2}{1-\rho^2} = \frac{0.2}{\xi^2}$	Variance of residual wages 0.2	Heathcote et al. (2010)
Taxes		
$T^y(y) = y - \delta y^{1-t_y}$, $t_y = 0.181$, $\delta = 0.9276$	Parametric approximation of the U.S. labor income tax schedule	Heathcote et al. (2017)
$T^b(b) = t_b b, \quad t_b = \begin{cases} 0.2, & b > 756{,}000 \\ 0, & \text{otherwise} \end{cases}$	Parametric approximation of the U.S. estate tax schedule	De Nardi and Yang (2016))
Education Cost		
$g(h', h) = \kappa \ln(h')^{\varsigma_1} \ln(h)^{\varsigma_2}$		
$\kappa = 10^{-8}$	Average years of schooling 12.8	Barro and Lee (2013)
$\varsigma_1 = 9.769$	Average net cost of one year of education 13,845	Stantcheva (2017)
$\varsigma_2 = -0.4399$	Intergenerational correlation of schooling 0.46	Hertz et al. (2008)

Table 2: Target Statistics in the Data and the Model

		Model			
Variable	**Data**	**Baseline**	$\boldsymbol{\vartheta = 0}$	$\boldsymbol{\vartheta = 0.05}$	**KP**
Average years of schooling S	12.86	12.52	12.50	13.43	12.67
Correlation (S', S)	0.46	0.47	0.41	0.73	0.46
Intergenerational earnings elasticity	0.45	0.44	0.44	0.46	0.44
Average net cost of an additional year of schooling	$13,845	$13,194	$13,227	$12,146	$14,424
Bequest–wealth ratio	0.008	0.007	0.008	-0.011	0.398
Education expenditures–income ratio	0.02	0.021	0.021	0.016	0.022

Notes: Statistics are computed conditional on dynasties whose parents have zero assets, 12 years of schooling, and ability persistence $\rho = 0.459$. The baseline model corresponds to the calibrated level of altruism used in the benchmark economy, while the alternative columns set $\vartheta = 0$ (no altruism) and $\vartheta = 0.05$ (high altruism).

Varying ϑ highlights the role of education-specific altruism in shaping intergenerational allocations. Higher values of ϑ increase education investment and average schooling. At the same time, stronger altruism shifts intergenerational transfers away from financial bequests toward human capital accumulation. For sufficiently high values of ϑ, the bequest–wealth ratio becomes negative, indicating that parents borrow to finance their children's education.

Stronger altruism also increases intergenerational persistence in schooling, as parents place greater weight on children's human capital and concentrate investments among already advantaged dynasties.

The education expenditure–income ratio plays a central role in disciplining the level of education-specific altruism. Because this moment responds monotonically to ϑ, it provides a key empirical anchor for the strength of parental preferences over children's human capital.

Calibration and identification. The parameter vector $(\kappa, \varsigma_1, \varsigma_2, \rho, \beta, \vartheta)$ is chosen to match a set of empirical moments reported in Table 2, targeting key features of education, intergenerational mobility, and wealth transmission.

The education cost parameters $(\kappa, \varsigma_1, \varsigma_2)$ are identified using moments on average schooling, the net cost of education, and the intergenerational correlation of schooling. The discount factor β is pinned down by the bequest–wealth ratio, while the persistence parameter ρ is disciplined by the intergenerational earnings elasticity.

The altruism parameter ϑ is calibrated to match the empirical moment that households spend approximately 2% of lifetime income on education. While this moment does not point-identify ϑ, the monotonic relationship between ϑ and education expenditures provides strong discipline. We therefore treat ϑ as a disciplined parameter and use variation in ϑ to quantify the importance of education-specific altruism in shaping optimal policy.

The model slightly understates the bequest–wealth ratio relative to the data. Given the small magnitude of this moment, the discrepancy is modest in absolute terms. The model slightly understates the bequest–wealth ratio (0.007 versus 0.008 in the data), a discrepancy of less than 13 percent. This reflects a mild tension between the bequest moment and the education expenditure moment: a higher β would close this gap but would push education expenditures above the 2% target. We prioritize the education expenditure moment given its direct relevance to the mechanisms studied in the paper. The implied discrepancy is quantitatively negligible and does not affect the main results, since the wedge decomposition in Section 4 is primarily driven by the transfer channel, which depends on the normalized bequest wedge $\tau_b^*/(1-\tau_b^*)$ rather than the level of bequests.

5.3 Optimal Investment in Children's Human Capital

Figure 1 shows how the constrained-efficient investment in children's human capital varies with parental income and bequests. Two robust patterns emerge. First, investment is increasing in parental labor income. This reflects the informational role of income: due to intergenerational persistence of ability, higher-income dynasties are more likely to have children with higher expected productivity, which raises the marginal return to education.

Second, investment is decreasing in bequests. Holding income fixed, larger bequests reduce the optimal level of human capital investment. This pattern reflects weaker labor supply incentives for children with higher inherited wealth, which lowers the social return to education. Taken together, these forces generate a pronounced cross-sectional gradient: investment is highest for dynasties with high income and low bequests.

The relationship is non-linear, with the sensitivity of investment to income declining at higher levels of inherited wealth, as wealth-induced distortions reduce the effectiveness of education in increasing future earnings.

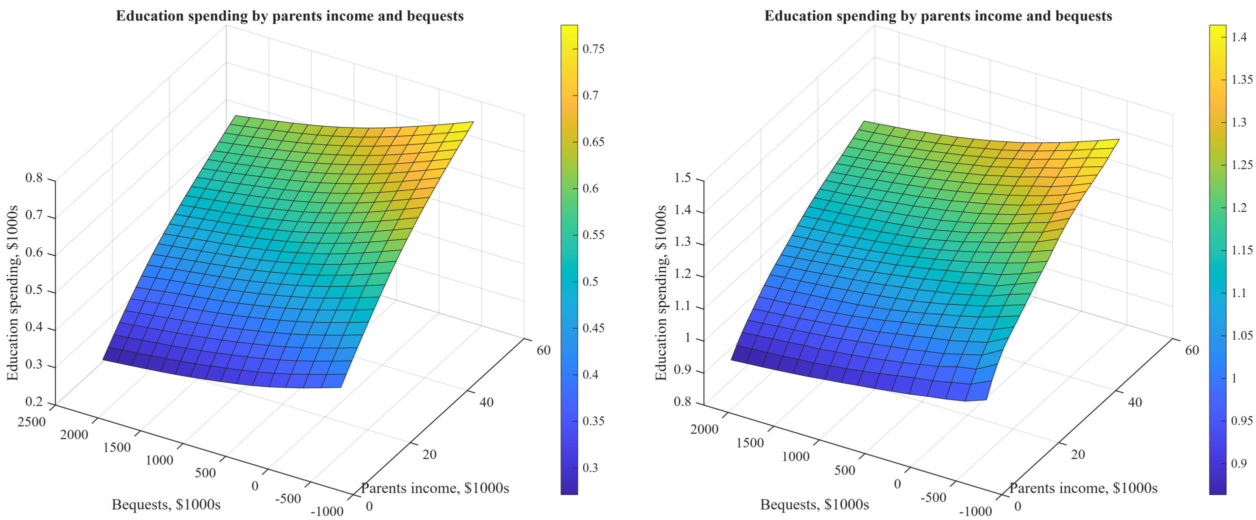

(a) Baseline

(b) High altruism ($\vartheta = 0.05$)

Figure 1: Optimal investment in children's human capital by parental income and bequests

Notes: The figure displays optimal investment in children's human capital as a function of parental labor income and inherited wealth (bequests). The left panel shows the baseline economy, while the right panel corresponds to an economy with high altruism ($\vartheta = 0.05$). We condition on parents with zero assets, 12 years of schooling, and persistence of ability equal to $\rho = 0.459$.

Comparing panels (a) and (b) reveals a systematic effect of parental altruism. Higher altruism shifts the entire investment schedule upward: for any given level of parental income and bequests, dynasties invest more in children's human capital when intergenerational linkages are stronger. At the same time, the sensitivity of investment to inherited wealth is attenuated. While investment remains decreasing in bequests, the gradient becomes flatter under high altruism, reflecting that stronger parental concern partially offsets wealth-induced distortions.

To understand the mechanisms underlying these patterns, we now turn to the constrained-efficient human capital wedge and its decomposition.

5.4 Human Capital and Bequest Wedges with Altruistic Parents

Koeniger and Prat (2018) establish a tight link between the wedges on bequests and human capital in dynastic models with asymmetric information. In their benchmark environment, the implicit tax on bequests is positive, whereas the human capital wedge is negative, reflecting underinvestment in education due to the risky nature of human capital accumulation.

The key contribution of this paper is to show that parental altruism can reverse the direction of optimal education policy: from subsidizing human capital investment, as in models without altruism, to taxing it.

Our framework separates two preference margins. The first is the dynastic discount factor β, which governs the private valuation of intergenerational transfers. The second is education-specific altruism ϑ, which captures the direct utility parents derive from their children's human capital. Relative to Koeniger and Prat (2018), both margins affect the human capital wedge: β through the intertemporal transfer motive, and ϑ through the private marginal valuation of education.

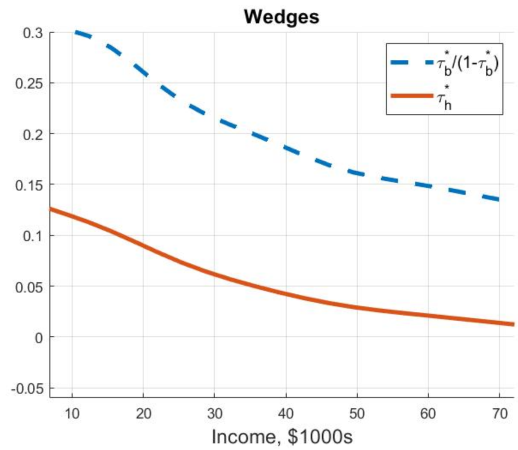

Notes: Baseline model with parental altruism ($\vartheta = 0.0005$).

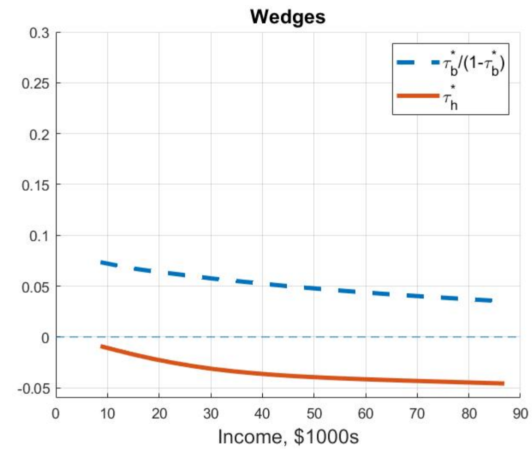

Notes: Benchmark economy without altruism, following Koeniger and Prat (2018).

Figure 2: Human capital wedges across generations

Figure 2 illustrates how the introduction of parental altruism modifies the structure of wedges. In the benchmark economy without altruism, the human capital wedge is negative, reflecting underinvestment in education. In contrast, in the baseline calibration with altruism, the wedge becomes positive, indicating that private incentives may exceed the social optimum.

To understand this pattern, it is useful to isolate the forces driving the wedge. The quantitative difference between the two panels is modest at the calibrated value $\vartheta = 0.0005$. A key feature of both economies is the large and positive bequest wedge $\tau_b^*/(1-\tau_b^*)$, driven by the gap between the dynastic discount factor $\beta = 0.17$ and the planner's discount factor $q = 0.412$. Through the transfer channel Δ_b, this intertemporal distortion propagates into the human capital wedge and shifts it toward positive values even at low levels of altruism. Education-specific altruism provides the marginal force that determines the sign of the wedge at moderate levels and becomes dominant only when ϑ is sufficiently large.

The difference between bequest and human capital wedges is driven by the risk associated with human capital investment. While bequests yield a relatively safe return, the payoff from education depends on the child's realized ability, which is uncertain at the time of investment. This risk makes human capital a poor hedge against consumption risk and limits the extent to which the planner discourages education relative to bequests.

Figure 3 decomposes the constrained-efficient human capital wedge into its two components, following Proposition 4: the intergenerational transfer component Δ_b and the incentive component Δ_i.

The decomposition shows that variation in the human capital wedge is primarily driven by the transfer channel. Differences in inherited wealth affect the marginal value of intergenerational transfers by altering the marginal utility of future consumption and labor income, which in turn changes the attractiveness of human capital investment as a vehicle for resource transmission across generations. As a result, the transfer component Δ_b accounts for most of the cross-sectional variation in the wedge, while the incentive component Δ_i plays a quantitatively smaller role.

Comparing across levels of altruism reveals that parental preferences also shape the underlying mechanism. Under high altruism, the human capital wedge increases substantially and becomes uniformly positive

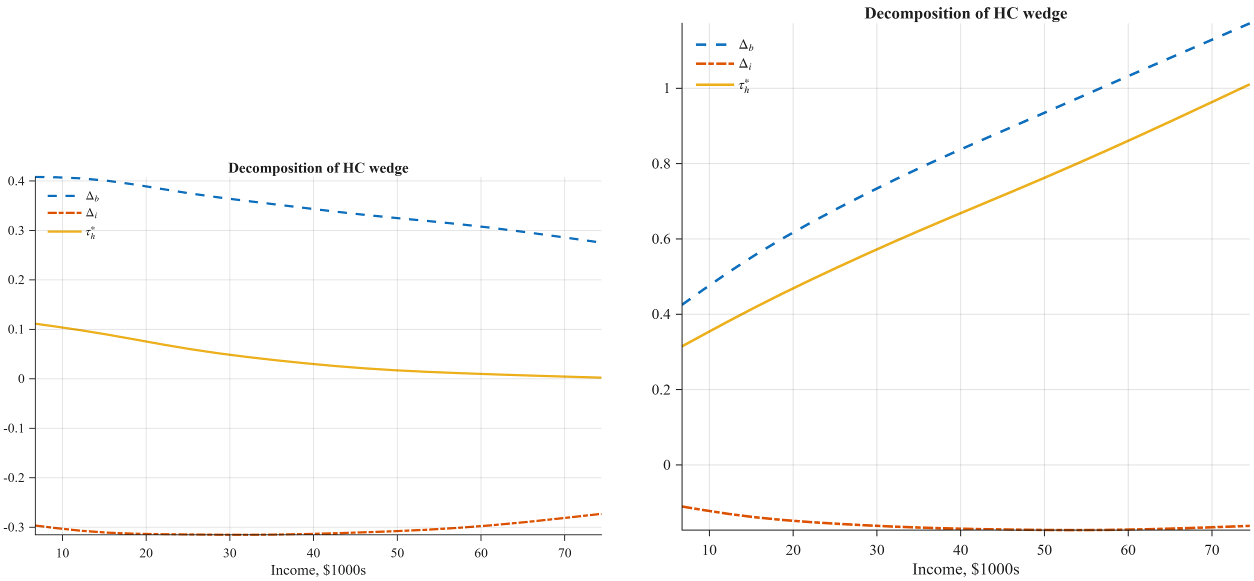

(a) Baseline

(b) High altruism

Figure 3: Decomposition of the human capital wedge by altruism

Notes: The figure decomposes the constrained-efficient human capital wedge into its intergenerational transfer component Δ_b and incentive component Δ_i. The left panel shows the baseline economy, while the right panel corresponds to the high-altruism economy.

across the income distribution. This change is driven by a strong amplification of the transfer component Δ_b, while the incentive component remains largely unchanged. Stronger intergenerational linkages increase the value of using human capital investment as a transfer device, leading the planner to introduce larger distortions to regulate these transfers.

These results imply that human capital investment acts as an implicit intergenerational transfer technology. As altruism increases, the planner's problem increasingly resembles one of regulating redistribution across generations rather than correcting informational frictions. Consequently, cross-sectional variation in education investment is governed primarily by transfer motives, and the decline in investment with wealth reflects endogenous reductions in the social return to education rather than binding constraints or misallocation.

The degree of parental altruism ϑ affects the magnitude of the human capital wedge and can, for sufficiently high values, reverse its sign, as we document above.

Under Cobb–Douglas technology, the incentive component vanishes and the human capital wedge reduces to the transfer component, $\tau_h^* = \Delta_b$, as characterized in Proposition 4. The covariance term in Δ_b is negative, reflecting the fact that human capital is a poor hedge against consumption risk. This further explains why the human capital wedge remains smaller than the bequest wedge.

Policy implications. Figure 4 compares human capital wedges across economies with different levels of parental altruism. In the absence of altruism (panel (b)), the human capital wedge is negative, reflecting underinvestment in education, consistent with Koeniger and Prat (2018). In the baseline calibration (panel (a)), the wedge is close to zero and slightly positive, indicating that parental altruism partially internalizes the social return to education. When altruism is sufficiently strong (panel (c)), the wedge becomes positive and increasing in income, implying overinvestment in education and justifying corrective taxation.

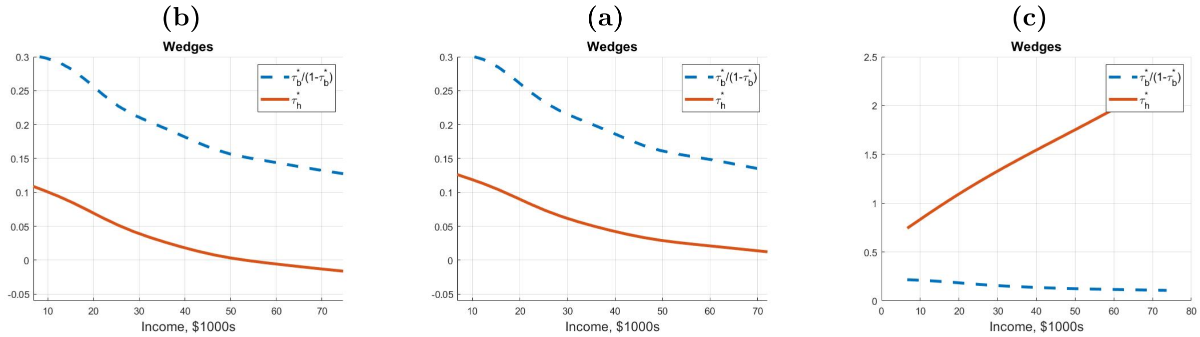

Notes: Economy without parental altruism ($\vartheta = 0$).

Notes: Baseline economy with moderate altruism ($\vartheta = 0.0005$).

Notes: Economy with high altruism ($\vartheta = 0.05$).

Figure 4: Human capital wedges under different levels of parental altruism

These patterns are consistent with the decomposition $\tau_h^* = \Delta_b + \Delta_i$: higher altruism modifies both the intergenerational transfer channel Δ_b and the incentive channel Δ_i through its effect on the equilibrium allocation.

Result 1 (Threshold education-specific altruism)**.** *Holding the remaining calibrated parameters fixed, there exists a threshold level of education-specific altruism such that the human capital wedge switches sign.*

The numerical result is consistent with continuity of the policy functions in ϑ. At low values of ϑ, the human capital wedge is negative or close to zero, whereas at high values it becomes positive, implying a reversal in the direction of optimal education policy.

Overall, the results highlight two distinct channels. The calibrated gap between β and the planner's discount factor governs the magnitude of wedges, while education-specific altruism ϑ can reverse their sign. When ϑ is sufficiently large, private incentives to invest in education exceed the social optimum, and the planner optimally taxes rather than subsidizes human capital investment.

6 Policy Experiments

6.1 Implementation of the Constrained-Efficient Allocation

We distinguish two levels of analysis. First, we characterize the full, history-dependent tax system that implements the constrained-efficient allocation. Second, Sections 6.2–6.3 evaluate simpler, administratively feasible instruments that approximate the optimum while keeping the existing U.S. income tax schedule fixed. These are second-best reforms, and their welfare gains should be interpreted relative to the status quo rather than to the constrained-efficient benchmark.

The constrained-efficient allocation can be decentralized through a history-dependent tax system $T_t(b_t, y_t, h_{t+1}, \theta^{t-1})$ that replicates the three wedges derived in Section 4. The marginal labor income tax implements the labor wedge directly, while the bequest tax is pinned down by the reciprocal Euler equation, which must hold ex post for each reported type. The marginal tax on human capital is linked to the human capital wedge through the household's optimality condition. This mapping is not one-to-one: a zero human capital wedge does not imply a zero marginal tax, and a positive wedge does not necessarily translate into a positive marginal tax (Koeniger and Prat, 2018; Stantcheva, 2020). Full derivations are provided in Appendix G.

In the following sections, we study two prominent education-financing instruments—income-contingent repayment loans and direct education subsidies—within the U.S. tax–transfer system.

6.2 Income-Contingent Repayment Loans

Income-contingent loan (ICL) programs, in which repayments are tied to future earnings, are widely used in countries such as the United States, the United Kingdom, and Australia. These systems expand access to higher education while providing insurance against income risk (Chapman et al., 2022). Empirical evidence from the United Kingdom suggests that such schemes do not generate significant distortions in labor supply (Britton and Gruber, 2020).

Implementation. In the model, families finance investments in children's human capital through education loans that are repaid by the next generation via history-dependent, income-contingent payments. Following Koeniger and Prat (2018), labor income taxes depend on contemporaneous earnings, while loan repayments depend on accumulated human capital and the family's financial history.

Labor income taxes follow the nonlinear specification

$$T_t^y(y_t) = y_t - \delta y_t^{1-t_y},$$

which approximates the U.S. income tax schedule. Parents may borrow to finance their children's education in the amount

$$L_t(h_{t+1}, h_t) = g(h_{t+1}, h_t),$$

while the next generation repays education debt through an income-contingent function

$$D_t\big(b_t, y_t, h_{t+1}, L^{t-1}, y^{t-1}\big)\,,$$

which depends on inherited wealth, realized income, and accumulated human capital.

The gross payment faced by households combines the existing income tax, an education loan, and an income-contingent repayment:

$$\underbrace{T_t(\cdot)}_{\text{gross payment}} = \underbrace{T_t^y(y_t)}_{\text{income tax}} - \underbrace{L_t(h_{t+1}, h_t)}_{\text{education loan}} + \underbrace{D_t(b_t, y_t, h_{t+1}, L^{t-1}, y^{t-1})}_{\text{loan repayment}}. \tag{44}$$

We focus on a simplified implementation in which education costs are fully financed through loans, while repayments depend only on realized income. Details of the household problem and derivations are provided in Appendix H.

Economic mechanism. ICLs relax borrowing constraints and provide insurance against income risk by linking repayments to realized earnings. As a result, they reduce the effective cost of human capital investment for low-income households while introducing an implicit tax on high-income states.

At the same time, income-contingent repayments act as an additional implicit tax on earnings in high-income states, introducing a countervailing distortion on labor supply that partially offsets the efficiency gains from improved insurance. The net welfare effect therefore reflects a trade-off between relaxed credit constraints and increased labor supply distortions at the top of the income distribution.

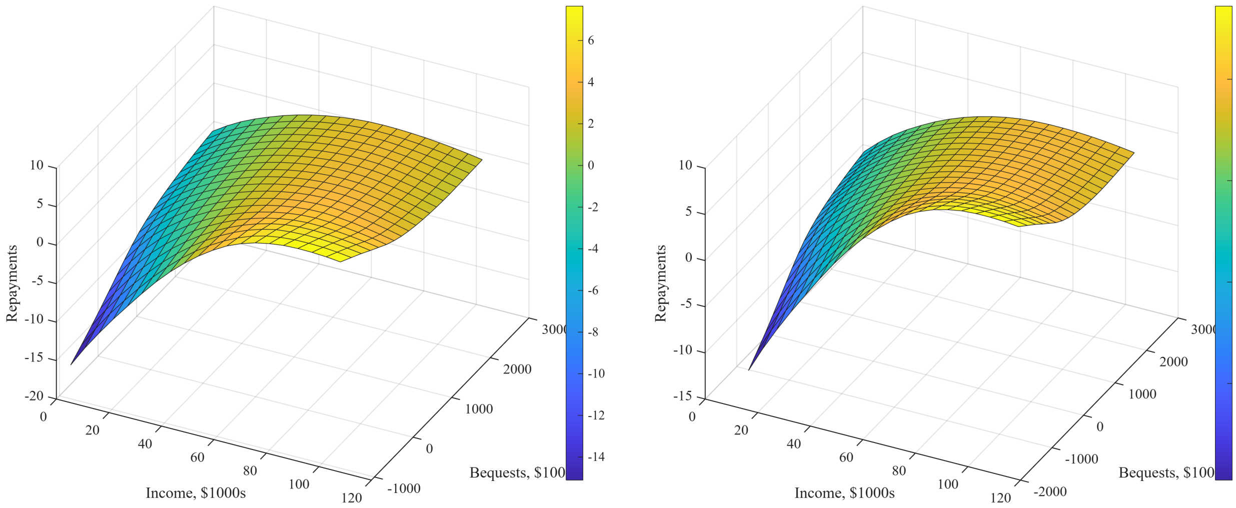

Baseline High altruism

Figure 5: Net payments under income-contingent loans.

Notes: Net payments are gross payments net of the education loan; see Section 6.2. Positive values indicate payments to the government, while negative values indicate net subsidies. The comparison shows how stronger parental altruism affects the distribution of net transfers.

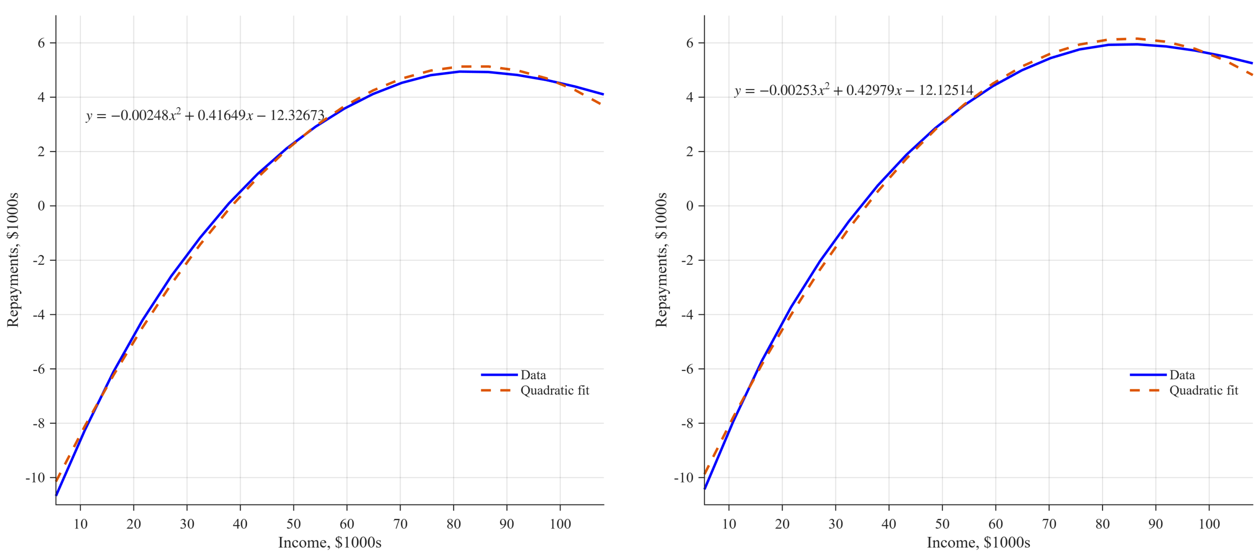

Baseline High altruism

Figure 6: Income-contingent repayment schedule.

Notes: In the simplified implementation, $L_t(h', h) = g(h', h)$ and education is fully loan-financed. Repayments depend on labor income: $D_t(y) = 0$ below a threshold and increase thereafter. The comparison highlights how stronger parental altruism modifies the repayment profile.

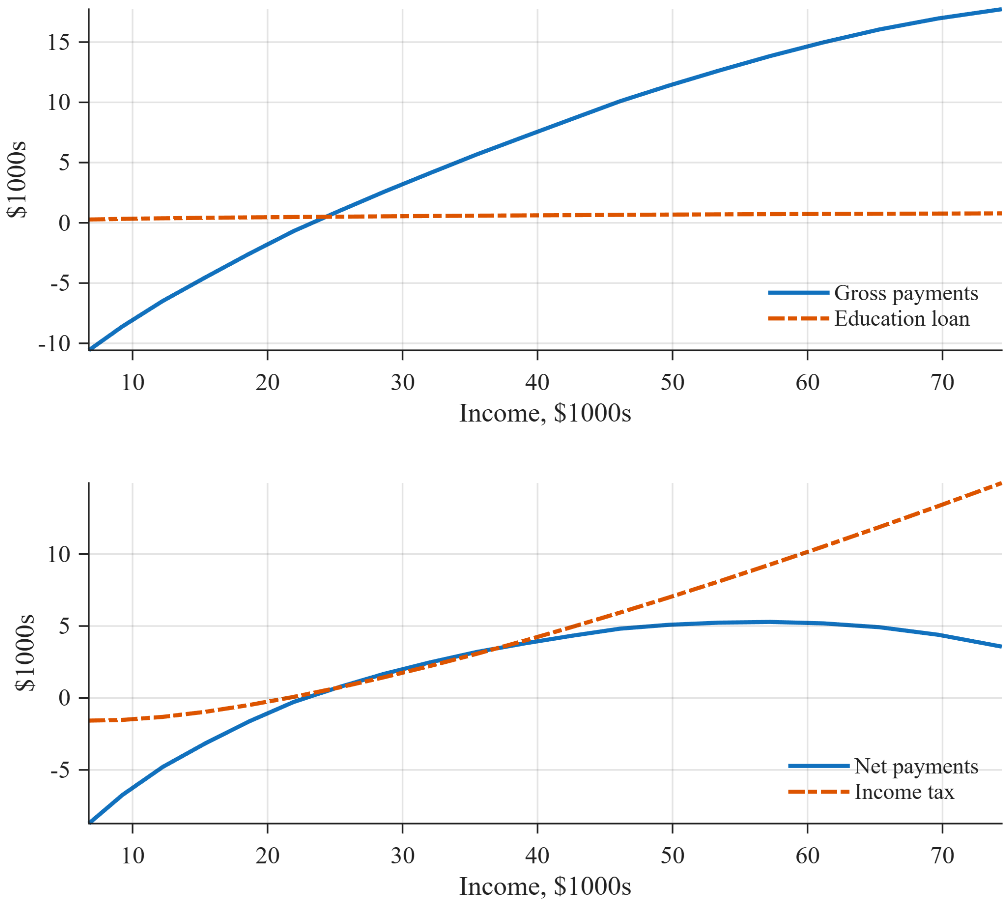

Figure 7: Implementation of the income-contingent loan system

Notes: The figure illustrates how labor income maps into payments under the income-contingent loan system. The top panel shows gross payments together with the education loan. The bottom panel reports net payments (gross payments net of the loan) alongside labor income taxes. At low income levels, individuals receive net transfers, while repayments increase with income beyond a threshold. The system therefore combines insurance against low earnings with progressive repayment.

Figure 7 illustrates this mechanism. At higher income levels, net payments increase more slowly, reflecting the trade-off between insurance and incentive provision in a progressive system. As a result, both low-income dynasties facing borrowing constraints and high-income dynasties with substantial bequests receive net education subsidies under the optimal policy. The implementation of the income-contingent loan system is largely unaffected by the degree of parental altruism. The repayment schedule and the distribution of net transfers remain nearly identical across calibrations. This contrasts with education subsidies, whose optimal design is highly sensitive to parental altruism, a comparison we explore in detail in Section 6.4.

Figure 5 shows net payments under the optimal ICL scheme as a joint function of labor income and bequests. Net repayments rise with income and decline sharply with bequests: households with large bequests receive net subsidies even at moderate income levels, reflecting the negative relationship between optimal policy and private resources. Under high altruism (right panel), the surface remains largely unchanged relative to the baseline. Differences across calibrations are quantitatively small, indicating that the structure of net payments is robust to the degree of parental altruism.

Figure 6 isolates the income gradient by conditioning on families with zero bequests. The repayment schedule is concave: repayments are negative at low incomes, rise through the middle of the distribution, and flatten above approximately \$60,000. This pattern reflects the insurance role of ICLs—subsidizing low earners who cannot repay—while limiting net transfers to high-income households facing lower income risk. Under high altruism, the repayment schedule is nearly identical to the baseline, with only minor quantitative differences. This confirms that the optimal repayment rule is largely invariant to the degree of parental altruism.

Role of parental altruism. While the marginal benefit of ICLs declines with ϑ as private intergenerational transfers partially substitute for public insurance, the structure of the optimal repayment schedule remains largely invariant to the degree of altruism. We assess the welfare implications of this asymmetry in Section 6.4.

Table 3: Quantitative Effects of Income-Contingent Loans

Variable	**Data**	**Status Quo**	**(With ICL**
Average years of schooling S	12.86	12.52	14.06
Correlation (S', S)	0.46	0.47	0.45
Intergenerational earnings elasticity	0.45	0.44	0.45
Average net cost of an additional year of schooling	\$13,845	\$13,194	\$11,783
Bequest–wealth ratio	0.008	0.007	0.004
Education expenditures–income ratio	0.01–0.03	0.02	0.02
Aggregate welfare		1.8161×10^5	3.0590×10^5
Aggregate output	\$$14.99 \times 10^{12}$	\$$1.2656 \times 10^{12}$	\$$1.5256 \times 10^{12}$
Aggregate income tax revenue (% of output)	10.67%	15.28%	18.33%

Notes: Model statistics are reported for $\vartheta = 0.0005$. Aggregate output in the data corresponds to U.S. GDP in 2010 (\$14.99 trillion, World Bank). The average individual income tax rate in the data equals 13.29 percent (Tax Foundation). Aggregate welfare is measured as the sum of utilities across all dynasties.

Quantitative results. ICLs lead to a substantial increase in human capital accumulation: average schooling rises by approximately 1.5 years relative to the status quo (Table 6), reflecting improved access to education for liquidity-constrained households and reduced downside risk associated with educational investment. This expansion translates into higher aggregate output and tax revenues, with income tax revenue increasing from 15.28% to 18.33% of output—a fiscal dividend that exceeds the direct cost of the program.

Table 4: Target Statistics in the Data and the Model with and without Income-Contingent Loans (ICL)

	Earnings Quantile	**Model $\vartheta = 0.0005$**	
		Status quo	**With ICL**
Years of schooling	0.25	12.37	13.81
	0.50	12.41	13.86
	0.75	12.51	13.93
	0.95	12.52	13.98
Cost of additional schooling	0.25	\$13,275	\$11,846
	0.50	\$13,247	\$11,825
	0.75	\$13,221	\$11,805
	0.95	\$13,200	\$11,787

Notes: Statistics are reported by earnings quantile for the benchmark model with $\vartheta = 0.0005$. The table compares the status quo education finance system with an economy featuring income-contingent repayment loans (ICL).

Table 4 reports the distributional effects across earnings quantiles. ICLs reduce disparities in education costs: the gap between low- and high-income households declines from \$75 to \$59. However, dispersion in years of schooling increases modestly, from 0.15 to 0.17 years. This pattern reflects heterogeneous responses across the income distribution. While access to education improves broadly, high-income and high-ability families—who face lower repayment risk and expect higher returns—expand educational investment more strongly. As a result, ICLs reduce financial inequality in education while modestly increasing inequality in educational attainment.

Overall, income-contingent loans improve efficiency by relaxing borrowing constraints and providing insurance against income risk. However, because they do not directly target heterogeneity in returns to education, their redistributive impact is limited — and their welfare gains are smallest in environments with strong intergenerational links, where private transfers already provide effective insurance.

6.3 Implementation of Education Subsidies

Public subsidies for higher education are widely used across countries, yet their effects on inequality and welfare remain theoretically ambiguous. Bovenberg and Jacobs (2005) and Stantcheva (2015) show that education subsidies and labor income taxes must be jointly designed to ensure efficient human capital investment. Extending this insight to an overlapping-generations environment with borrowing constraints, Stantcheva (2015) characterizes optimal education subsidies as functions of bequests and labor income taxes.

In this section, we complement the income-contingent loan analysis of Section 6.2 by examining a second prominent policy instrument: direct education subsidies. The government finances a fraction τ_g of education expenditures $g(h', h)$, while families pay the remaining share $(1 - \tau_g)g(h', h)$.

We hold the labor income tax schedule fixed and evaluate fiscal feasibility ex post by comparing tax revenues with subsidy expenditures. We normalize $\tau_g = 0$ to represent the current U.S. system, so that the calibrated cost function already incorporates existing public support. Positive values of τ_g therefore correspond to additional subsidies relative to this baseline.

Under education subsidies, the family's dynamic optimization problem is given by

$$\hat{W}_t(b_t, h_t, \theta_t) = \max_{b_{t+1}, h_{t+1}, l_t} \left\{ U(c_t, l_t, h_{t+1}) + \beta \int_{\Theta} \hat{W}_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1}) \, dF(\theta_{t+1}|\theta_t) \right\}, \tag{45}$$

subject to the budget constraint

$$b_{t+1} + c_t + (1-\tau_g) g(h_{t+1}, h_t) = (1+r) b_t + y(l_t, h_t, \theta_t) - T(b_t, y(l_t, h_t, \theta_t), \theta_t)\,. \tag{46}$$

Relative to the implementation discussed in Section 6.1, this formulation differs only by the introduction of an explicit education subsidy τ_g and by assuming that human capital investments are not directly taxed, i.e. $T(h_{t+1}) = 0$. While education subsidies are typically financed through general taxation, we abstract from the government budget constraint and focus on a partial-equilibrium environment. The objective is to isolate the effects of education subsidies on household decisions regarding consumption, labor supply, and investment in children's human capital.

The marginal income tax, the labor wedge, and the bequest wedge are unchanged. The only difference concerns the relationship between the marginal tax on human capital and the human capital wedge. Using human capital wedge from (1) together with the envelope condition, the first-order condition for human capital can be written as

$$(1-\tau_g)\left(\frac{\partial g(h_{t+1},h_t)}{\partial h_{t+1}} + \beta \mathbb{E}_t\left[\frac{\frac{\partial U(c_{t+1},l_{t+1},h_{t+2})}{\partial c_{t+1}}}{\frac{\partial U(c_t,l_t,h_{t+1})}{\partial c_t}} \frac{\partial g(h_{t+2},h_{t+1})}{\partial h_{t+1}}\right]\right) = \frac{\frac{\partial U(c_t,l_t,h_{t+1})}{\partial h_{t+1}}}{\frac{\partial U(c_t,l_t,h_{t+1})}{\partial c_t}} + \beta \mathbb{E}_t\left[\frac{\frac{\partial U(c_{t+1},l_{t+1},h_{t+2})}{\partial c_{t+1}}}{\frac{\partial U(c_t,l_t,h_{t+1})}{\partial c_t}} \frac{\partial y_{t+1}}{\partial h_{t+1}}\left(1 - \frac{\partial T_{t+1}}{\partial y_{t+1}}\right)\right]. \tag{47}$$

Equation (47) equates the marginal cost and the marginal benefit of investing in human capital. This condition highlights the interdependence between education subsidies and labor income taxation in shaping optimal human capital investment. Parental altruism increases the perceived benefit from education, allowing the optimality condition to be satisfied with either higher labor taxes or lower education subsidies relative to models without altruism.

We back out the education subsidy τ_g that decentralizes the constrained-efficient allocation, conditional on family characteristics. This subsidy reflects the marginal distortion in human capital investment implied by the planner's allocation. When expected returns to education are high, the optimal policy may involve a negative subsidy, that is, a tax on human capital investment, in order to mitigate inequality arising from heterogeneous returns.

However, because returns to education are uncertain ex ante and not directly observable by the policymaker, such fully state-contingent subsidies are not implementable in practice. We therefore use these ex-post optimal subsidies as a benchmark to evaluate feasible policy designs.

To assess implementable policies, we conduct two exercises. First, we compute the ex-post optimal subsidy for each family type based on realized characteristics. Second, we consider two implementable schemes—a uniform subsidy and an income-contingent subsidy—and quantify their effects on education investment and welfare.

Figure 8 shows how the optimal education subsidy τ_g varies jointly with parental labor income and bequests. Both panels reveal a highly heterogeneous subsidy schedule, but differ markedly in level, range, and in the relative importance of the two state variables.

In the baseline calibration (left panel), the subsidy ranges from approximately -1 to 0.8. Low-income households with modest bequests receive the largest positive transfers, while high-income households with

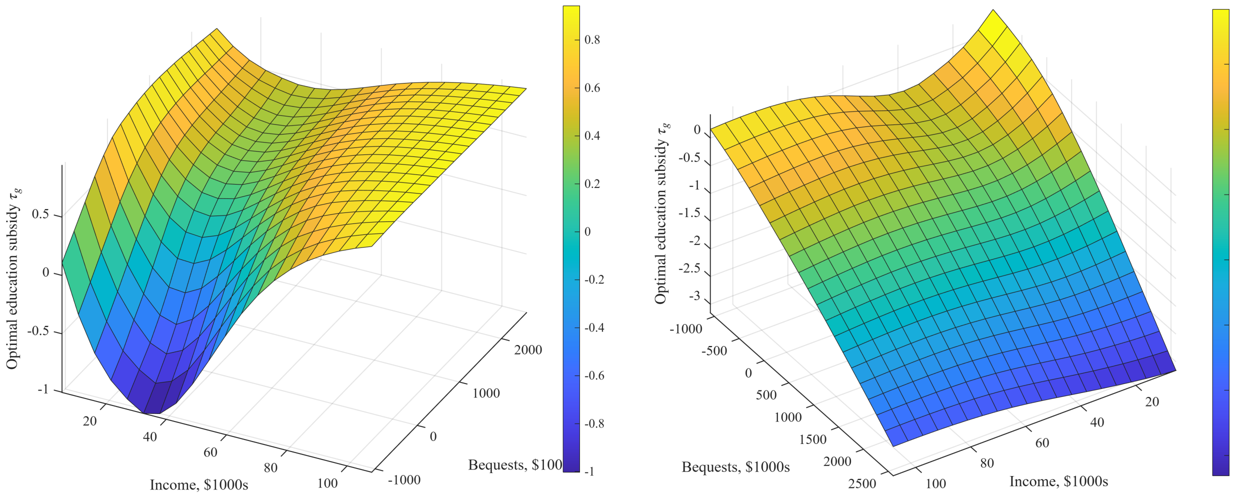

Baseline High altruism

Figure 8: Optimal education subsidies as a function of labor income and bequests

Notes: The figures show the optimal education subsidy τ_g implied by the constrained-efficient allocation. The left panel corresponds to the baseline calibration, while the right panel shows the case of extremely high altruism. Positive values correspond to subsidies, while negative values indicate taxes on human capital investment. The subsidy varies with parental labor income and bequests, reflecting a trade-off between redistribution and incentive alignment in the presence of heterogeneous returns to education.

large bequests optimally face taxes on human capital investment. The dependence on bequests is present but moderate: at a given income level, higher bequests reduce the subsidy, but variation across income remains the primary driver of the overall shape.

The high-altruism panel (right panel) reveals a qualitative shift in the optimal policy. The entire subsidy schedule moves into negative territory, with τ_g ranging from approximately -3 to 0, so that no household receives a positive education subsidy. In addition, the role of bequests becomes substantially stronger: at a given income level, the subsidy declines steeply with parental wealth, while the income gradient becomes noticeably flatter.

This pattern reflects two reinforcing mechanisms. First, when parents are highly altruistic, they internalize a large share of the social return to education, reducing the need for corrective public transfers. Second, wealthy altruistic dynasties tend to overinvest in education from a social perspective, shifting the role of government from correcting underinvestment to restraining excessive investment.

Taken together, these forces imply that bequests becomes the key state variable for policy in high-altruism environments, and that optimal policy takes the form of a tax on human capital investment whose magnitude increases sharply with bequests.

Figure 9 complements the previous figure by reporting the optimal subsidy schedule τ_g as a function of labor income alone, conditioning on families with zero bequests. This cross-section isolates the income gradient of the policy and allows a direct comparison between the two calibrations.

In the baseline (left panel), the subsidy schedule follows a pronounced U-shape. Low-income households receive substantial positive transfers, while the subsidy becomes negative at middle incomes before rising again for high-income households. This non-monotonic pattern reflects the trade-off between redistribution toward low earners and incentive alignment for higher-income households with stronger returns to education. The quadratic fit captures the overall curvature of the schedule, although it smooths over the sharp trough

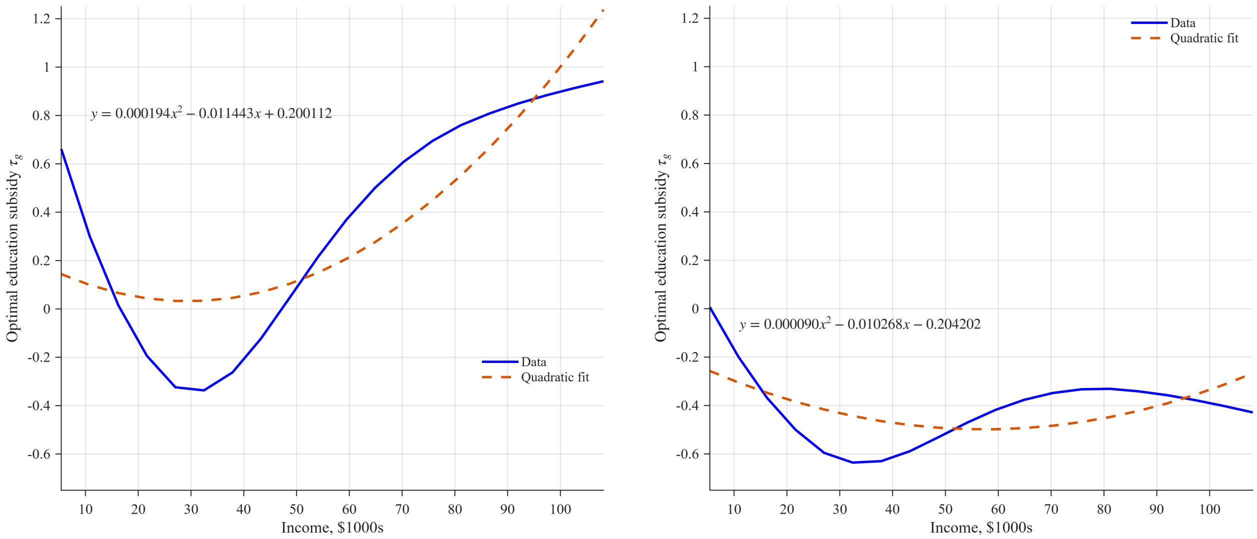

Baseline High altruism

Figure 9: Optimal education subsidies as a function of labor income (zero bequests)

Notes: The figures show the optimal education subsidy τ_g implied by the constrained-efficient allocation as a function of labor income for families with zero bequests. The left panel corresponds to the baseline calibration, while the right panel shows the case of extremely high altruism. The solid line represents the computed optimal subsidy, while the dashed line shows a quadratic approximation used for policy implementation. Positive values correspond to subsidies, while negative values indicate taxes on human capital investment. The comparison highlights that stronger parental altruism shifts the subsidy schedule downward, reflecting reduced need for public support due to stronger private incentives to invest in human capital.

at middle incomes.

In the high-altruism case (right panel), the U-shape is preserved, but the entire schedule shifts downward so that subsidies become negative across the entire income distribution. While the curvature remains similar, the level of the policy is uniformly lower, reflecting the reduced need for public intervention.

The comparison across panels highlights that parental altruism does not alter the qualitative structure of the optimal income-dependent policy, but significantly lowers its level. In effect, altruism acts as a substitute for public education subsidies: as families internalize a larger share of the returns to human capital investment, the planner optimally reduces or even reverses transfers.

To assess the welfare effects of implementing the optimal education subsidy, we use the interpolated subsidy rule

$$\tau_g(y) = 0.000194y^2 - 0.011443y + 0.200112,$$

where income y corresponds to labor output in the model. Computational details are provided in Appendix I.

We compare two policy implementations: a uniform subsidy $\tau_g = 0.35$ applied equally to all families, and the income-dependent rule $\tau_g(y)$ defined above. In both cases, families finance the remaining share $(1 - \tau_g)g(h', h)$ of education costs.

Under the uniform subsidy, all families receive identical support regardless of income or assets. This policy increases average years of schooling from 12.52 to 12.77 (Table 5), as parents reallocate resources from bequests toward education. Aggregate output and aggregate welfare also increase. Importantly, the uniform subsidy has only a minimal effect on labor supply. Since the subsidy is not conditioned on income, it does not significantly distort work incentives.

Table 5: Target Statistics by Earnings Quantile under Alternative Education Subsidy Schemes

	Earnings Quantile	**Model $\vartheta = 0.0005$**		
		$\tau_g = 0.0$	$\tau_g = 0.35$	$\tau_g(y)$
Years of schooling	0.25	12.37	12.61	12.80
	0.50	12.41	12.66	12.81
	0.75	12.51	12.71	12.83
	0.95	12.52	12.76	12.87
Cost of additional schooling	0.25	$13,275	$8,439	$6,528
	0.50	$13,247	$8,421	$6,996
	0.75	$13,221	$8,403	$7,316
	0.95	$13,200	$8,389	$7,420

Notes: Statistics are reported by earnings quantile for the benchmark model with $\vartheta = 0.0005$. Columns compare the baseline economy without education subsidies ($\tau_g = 0$), a uniform subsidy ($\tau_g = 0.35$), and an income-dependent subsidy given by $\tau_g(y) = 0.000194y^2 - 0.011443y + 0.200112$.

Higher educational attainment raises labor productivity and, consequently, tax revenues. For example, when $\tau_g = 0.35$, labor income tax revenue increases by $11.75 billion, exceeding the total cost of education subsidies of $9.03 billion. Hence, the uniform subsidy is self-financing and welfare-improving.

Under the income-dependent subsidy, low- and middle-income families receive more generous public support. This policy structure raises educational attainment but introduces distortions in labor supply, as households may reduce labor effort in order to qualify for higher subsidies. As a result, aggregate output and tax revenues decline (Table 5). Nevertheless, the government budget remains balanced when income tax revenues are allocated exclusively to education expenditures.

We next assess the distributional effects of education subsidies by comparing years of schooling and net education costs across earnings quantiles. This comparison illustrates how human capital investment varies along the income distribution and whether different subsidy regimes mitigate educational disparities.

Under the uniform subsidy regime, all families receive the same level of public support, independent of income or wealth. While this policy lowers the absolute cost of education for all households, it has only a limited effect on educational inequality. The gap in average years of schooling between low- and high-income families remains largely unchanged relative to the baseline without subsidies. Although uniform subsidies reduce differences in gross education costs, low-income households continue to face higher net costs due to their lower initial human capital, consistent with the implications of the Ben-Porath (1967) human capital production function.

In contrast, income-dependent subsidies generate a different distributional pattern. Because subsidy rates are higher for lower-income families, these households face a lower marginal cost of education. This strengthens incentives to invest in human capital among disadvantaged groups and leads to a reduction in the schooling gap between low- and high-income families. As illustrated in Figure 9, the subsidy schedule is explicitly designed to tilt marginal incentives toward low earners.

The distinct structures of the two subsidy regimes also have important efficiency implications. Uniform subsidies affect the marginal cost–benefit condition for all families symmetrically, thereby preserving labor supply incentives across the income distribution. Income-dependent subsidies, by contrast, introduce heterogeneity in marginal incentives and distort labor supply decisions, particularly among middle-income households seeking to qualify for higher support.

Role of parental altruism. The welfare gains from education subsidies depend on the degree of parental altruism ϑ. When altruism is low, families underinvest in education relative to the social optimum, and subsidies generate large welfare gains by correcting this distortion. As altruism increases, families internalize a larger share of the social return to education, reducing the corrective role of public subsidies. For sufficiently high values of ϑ, families may overinvest, and the optimal policy involves taxing rather than subsidizing education. The interaction between altruism and education subsidies is assessed systematically in Section 6.4.

In summary, income-dependent subsidies reduce inequality in educational attainment but do so at the cost of lower aggregate output and welfare due to behavioral distortions. Uniform subsidies, in contrast, increase aggregate welfare and output while leaving the distribution of educational attainment largely unchanged.

Our analysis is conducted in a partial-equilibrium environment, in which wages are fixed and fiscal feasibility is evaluated by comparing income tax revenues with education expenditures. An important extension would embed the model in general equilibrium, allowing wages, labor taxes, and education subsidies to be jointly determined. Such an extension would enable a richer analysis of feedback effects on wage dispersion, returns to education, and long-run fiscal sustainability.

6.4 Comparison of Income-Contingent Loans and Education Subsidies

This section compares the optimal design and macroeconomic implications of income-contingent loans (ICLs) and education subsidies (ES), focusing on their effects on human capital accumulation, aggregate output, inequality, and welfare.

A common implication emerging from the preceding analysis is that both policy instruments should be negatively related to parental bequests. Wealthy families already possess the means to finance education privately and face weaker marginal incentives to supply labor; additional public support to such households is therefore inefficient. Moreover, both ICLs and ES display a non-monotonic, parabolic relationship with labor income: households at the lower and upper ends of the income distribution receive the largest support, while middle-income families benefit comparatively less. This pattern mirrors the structure of the U.S. tax and transfer system, where low-income households are credit constrained, high-income households face high marginal taxes, and middle-income households typically have sufficient resources to invest in education without substantial government intervention. Nevertheless, the presence of a human capital wedge implies that policy intervention remains welfare-improving even for these middle-income groups.

Despite these similarities in optimal targeting, income-contingent loans and education subsidies operate through distinct mechanisms and therefore generate different macroeconomic and distributional outcomes. Income-contingent loans primarily relax borrowing constraints and insure households against income risk. By allowing education costs to be repaid conditionally on realized earnings, ICLs encourage educational investment across the income distribution. In the model, this mechanism raises average educational attainment, aggregate output, and tax revenues. However, these gains come at the cost of increased educational inequality: while college attendance rises for both low- and high-income households, the relative expansion is larger for high-income and high-ability families, leading to a widening gap in human capital.

Education subsidies, by contrast, directly reduce the marginal cost of education and thus affect investment incentives more immediately. Their effects crucially depend on whether subsidies are uniform or income dependent. Uniform subsidies increase educational attainment and aggregate output without substantially altering the distribution of education. Because the subsidy does not vary with income, labor supply incentives are largely preserved, and the resulting productivity gains translate into higher output and welfare. In the presence of parental altruism, these gains are amplified, as parents derive direct utility from their children's human capital.

Income-dependent education subsidies, however, introduce an explicit redistributive component. By providing larger subsidies to low-income households, they reduce educational inequality and compress the distribution of human capital. At the same time, this design distorts labor supply incentives, particularly among low- and middle-income households that adjust labor effort to qualify for higher subsidies. As shown in Figure 9, these distortions reduce aggregate output and tax revenues. The underlying mechanism is standard: publicly financed education partially crowds out private investment and labor supply, lowering taxable income. Importantly, this efficiency loss does not arise under uniform subsidies or under ICLs, where repayments are tied to future income rather than current labor supply decisions. A key asymmetry emerges from the comparison: parental altruism ϑ has a large and qualitatively transformative effect on optimal education subsidies, but leaves the structure of income-contingent loans largely unchanged. This asymmetry reflects the fact that the two instruments address fundamentally different market failures.

Education subsidies correct the gap between private and social returns to human capital investment. Because altruistic parents already internalize part of this social return through direct utility from their children's education, higher ϑ reduces the wedge between private and constrained-efficient incentives, lowering the need for public intervention and, at sufficiently high levels, reversing its sign entirely. In this sense, parental altruism acts as a close substitute for education subsidies: both operate on the same margin—the marginal cost of human capital investment—and higher altruism crowds out the public instrument.

Income-contingent loans, by contrast, primarily address borrowing constraints and income risk. These frictions are orthogonal to ϑ: altruistic parents may invest more in their children's education, but they cannot eliminate the credit constraints that prevent low-income households from financing education upfront, nor can they insure their children against adverse income realizations. The insurance and credit-relaxation functions of income-contingent loans therefore remain valuable regardless of the degree of parental altruism, and the optimal repayment schedule is comparatively stable across calibrations. Unlike education subsidies, which are highly sensitive to parental altruism, the design of income-contingent loans is largely invariant to intergenerational linkages.

This distinction has a direct policy implication: in economies where parental altruism is strong, the case for education subsidies is substantially weakened, while the case for income-contingent loans remains robust.

Overall, the comparison highlights a fundamental equity–efficiency trade-off. Income-contingent loans are more effective at increasing aggregate output and welfare by promoting human capital accumulation with minimal distortions to labor supply. Education subsidies, especially when income dependent, are more effective at reducing inequality in educational attainment and human capital but do so at the cost of lower aggregate output.

Which instrument is preferred therefore depends on the policymaker's objective. If the primary goal is to maximize aggregate output and welfare, income-contingent loans dominate. If the objective is to reduce inequality in education and intergenerational mobility, income-dependent education subsidies are more effective. In the calibrated model with altruistic parents and a U.S.-style tax system, both instruments generate welfare gains relative to the status quo, but they do so through distinct channels and with different distributional consequences.

6.5 Policy Discussion and External Validity

The analysis highlights that the optimal design of education finance depends critically on how different policy instruments interact with parental resources, incentives, and labor supply responses. While both income-contingent loans and education subsidies can improve welfare, they operate through distinct channels and therefore imply different policy trade-offs.

Income-contingent loans primarily relax borrowing constraints and provide insurance against income risk without distorting contemporaneous labor supply. As a result, they are particularly effective at increasing aggregate human capital, output, and welfare. Education subsidies, by contrast, operate directly on the marginal cost of investment and can be targeted toward disadvantaged households. This makes them more effective at reducing educational inequality, but at the cost of stronger labor-supply distortions when conditioned on current income.

These results are consistent with observed policy designs. Income-contingent loan systems implemented in countries such as Australia, the United Kingdom, and the United States combine broad access to higher education with repayment schemes tied to future earnings, thereby limiting ex ante distortions. At the same time, many countries rely on targeted education subsidies to promote equity, despite concerns about behavioral responses.

The model provides a unified framework for understanding these policy choices. It highlights that no single instrument dominates across all dimensions: income-contingent loans perform well on efficiency grounds, while targeted subsidies are more effective for redistribution. The optimal mix of instruments therefore reflects a fundamental equity–efficiency trade-off and depends on societal preferences over redistribution, growth, and intergenerational mobility.

6.6 Limitations

Several simplifying assumptions limit the scope of the analysis and point to promising directions for future research.

First, the model is solved in partial equilibrium, with wages and the aggregate production structure taken as given. The substantial increase in human capital induced by income-contingent loans would, in general equilibrium, affect wage premia and the tax base, with quantitatively uncertain net effects.

Second, the analysis assumes that policy instruments can be conditioned on income and, in some cases, on parental bequests. In practice, wealth is imperfectly observed and subject to avoidance, which may constrain the feasibility of strongly targeted policies.

Third, the model assumes a homogeneous degree of parental altruism across dynasties. Allowing for heterogeneity in altruism could generate richer patterns of education investment and may strengthen the case for means-tested or multidimensional policy instruments.

Fourth, the quantitative results focus on steady-state comparisons and abstract from transition dynamics. Policy reforms such as the introduction of income-contingent loans may generate redistribution across cohorts during the transition, which is not captured in the steady-state analysis.

Finally, fertility is exogenously fixed. Allowing for endogenous family size would permit an analysis of how parents trade off investment per child and fertility decisions in response to education policies.

Addressing these extensions would provide a more complete assessment of the interaction between parental altruism, education investment, and optimal policy design.

7 Conclusion

This paper studies optimal education finance in a dynastic Mirrlees framework with heterogeneous families and education-specific parental altruism. When parents derive direct utility from their children's human capital, education serves not only as an investment in future earnings but also as a vehicle for intergenerational transfers. This additional preference channel fundamentally alters the design of optimal education policy under asymmetric information.

The central result is that parental altruism changes the nature of the human-capital wedge. By increasing the private marginal value of education, altruism reduces the gap between private and social incentives for human capital accumulation and can, for sufficiently strong preferences, reverse its sign. In contrast to standard models with selfish agents—where underinvestment is pervasive and subsidies are uniformly optimal—altruistic dynasties may overinvest in education, implying that corrective taxation rather than subsidization becomes optimal for some groups.

This mechanism also generates a sharp distinction between policy instruments. Education subsidies operate directly on the same margin as parental altruism and are therefore highly sensitive to it: as altruism increases, the case for subsidies weakens and may reverse. Income-contingent loans, by contrast, primarily address borrowing constraints and income risk, and are therefore largely robust to variation in altruistic preferences.

Quantitatively, the model implies that optimal education support is non-monotonic across the income and wealth distributions, reflecting the interaction between private resources, incentive constraints, and intergenerational transfers. Income-contingent loans improve efficiency by expanding access and providing insurance, while education subsidies are more effective at reducing inequality but at the cost of stronger labor-supply distortions. The relative desirability of these instruments therefore depends on the strength of parental altruism and the policy objective.

Overall, the results highlight that education policy cannot be designed independently of family preferences and the tax system. Accounting for parental altruism is essential for understanding both the level and the targeting of optimal education support, and helps reconcile why observed education policies differ across countries and income groups.

Several extensions would be useful for future work. Allowing for heterogeneity in parental altruism could generate richer patterns of education investment and policy targeting across households. Incorporating general equilibrium effects would capture how large changes in human capital affect wages and tax revenues. Finally, introducing transition dynamics would allow the analysis of reform paths and intergenerational redistribution during policy changes.

References

Albanesi, Stefania and Christopher Sleet (2006) "Dynamic optimal taxation with private information," *The Review of Economic Studies*, 73 (1), 1–30.

Attanasio, Orazio P. and Katja M. Kaufmann (2014a) "Education Choices and Returns to Schooling: Mothers' and Youths' Subjective Expectations and Their Role by Gender," *Journal of Development Economics*, 109, 203–216, 10.1016/j.jdeveco.2014.04.003.

Attanasio, Orazio P and Katja M Kaufmann (2014b) "Education choices and returns to schooling: Mothers' and youths' subjective expectations and their role by gender," *Journal of Development Economics*, 109, 203–216.

Barro, Robert J and Jong Wha Lee (2013) "A new data set of educational attainment in the world, 1950–2010," *Journal of development economics*, 104, 184–198.

Becker, Gary S. and Nigel Tomes (1979) "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility," *Journal of Political Economy*, 87 (6), 1153–1189, 10.1086/260831, Seminal article introducing a dynastic model where altruistic parents allocate income between consumption and investment in children's human and nonhuman capital, without considering financial bequests or asymmetric information.

——— (1986) "Human Capital and the Rise and Fall of Families," *Journal of Labor Economics*, 4 (3), S1–S39, 10.1086/298118, Extension of the Becker–Tomes framework analysing how family investments in human capital affect income distribution across generations.

Bernard, Tanguy, Stefan Dercon, Kate Orkin, and Alemayehu Seyoum Taffesse (2019) "Parental Aspirations for Children's Education: Is There a "Girl Effect"? Experimental Evidence from Rural Ethiopia," *AEA Papers and Proceedings*, 109, 127–132, 10.1257/pandp.20191072.

Bovenberg, A Lans and Bas Jacobs (2005) "Redistribution and education subsidies are Siamese twins," *Journal of Public Economics*, 89 (11-12).

Britton, Jack and Jonathan Gruber (2020) "Do income contingent student loans reduce labor supply?" *Economics of Education Review*, 79, 102061.

Card, David (1999) "The causal effect of education on earnings," *Handbook of labor economics*, 3, 1801–1863.

Carroll, Christopher D (2006) "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics letters*, 91 (3), 312–320.

Caucutt, Elizabeth M and Lance Lochner (2020) "Early and late human capital investments, borrowing constraints, and the family," *Journal of Political Economy*, 128 (3), 1065–1147.

Chapman, Bruce (2006) "Income contingent loans for higher education: International reforms," *Handbook of the Economics of Education*, 2, 1435–1503.

Chapman, Bruce Dearden et al. (2022) "Income-contingent loans in higher education financing," *IZA World of Labor*.

Chetty, Raj (2012) "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply," *Econometrica*, 80 (3), 969–1018.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014) "Where is the Land of Opportunity," *The geography of*.

Chevalier, Arnaud (2004) "Parental Education and Child's Education: A Natural Experiment," IZA Discussion Paper 1153, IZA Institute of Labor Economics.

Cunha, Flavio and James Heckman (2007) "The technology of skill formation," *American economic review*, 97 (2), 31–47.

De Nardi, Mariacristina and Fang Yang (2016) "Wealth inequality, family background, and estate taxation," *Journal of Monetary Economics*, 77, 130–145.

Doepke, Matthias and Fabrizio Zilibotti (2017) "Parenting with style: Altruism and paternalism in intergenerational preference transmission," *Econometrica*, 85 (5), 1331–1371.

Farhi, Emmanuel and Iván Werning (2013) "Insurance and taxation over the life cycle," *Review of Economic Studies*, 80 (2), 596–635.

Fernandes, Ana and Christopher Phelan (2000) "A recursive formulation for repeated agency with history dependence," *Journal of Economic Theory*, 91 (2), 223–247.

Galor, Oded and Joseph Zeira (1993) "Income distribution and macroeconomics," *The review of economic studies*, 60 (1), 35–52.

Giannola, Michele (2024) "Parental Investments and Intra-household Inequality in Child Human Capital: Evidence from a Survey Experiment," *The Economic Journal*, 134 (658), 671–727, 10.1093/ej/uead080.

Guryan, Jonathan, Erik Hurst, and Melissa Kearney (2008) "Parental education and parental time with children," *Journal of Economic perspectives*, 22 (3), 23–46.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante (2010) "The macroeconomic implications of rising wage inequality in the United States," *Journal of political economy*, 118 (4), 681–722.

——— (2017) "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132 (4), 1693–1754.

Hertz, Tom, Tamara Jayasundera, Patrizio Piraino, Sibel Selcuk, Nicole Smith, and Alina Verashchagina (2008) "The inheritance of educational inequality: International comparisons and fifty-year trends," *The BE Journal of Economic Analysis & Policy*, 7 (2).

Hintermaier, Thomas and Winfried Koeniger (2010) "The method of endogenous gridpoints with occasionally binding constraints among endogenous variables," *Journal of Economic Dynamics and Control*, 34 (10), 2074–2088.

Hotz, V Joseph, Emily E Wiemers, Joshua Rasmussen, and Kate Maxwell Koegel (2023) "The role of parental wealth and income in financing children's college attendance and its consequences," *Journal of Human Resources*, 58 (6), 1850–1880.

Kapička, Marek (2013) "Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach," *Review of Economic Studies*, 80 (3), 1027–1054.

Koeniger, Winfried and Julien Prat (2018) "Human capital and optimal redistribution," *Review of Economic Dynamics*, 27, 1–26.

Loury, Glenn C (1981) "Intergenerational transfers and the distribution of earnings," *Econometrica: Journal of the Econometric Society*, 843–867.

Milgrom, Paul and Ilya Segal (2002) "Envelope theorems for arbitrary choice sets," *Econometrica*, 70 (2), 583–601.

Milgrom, Paul and Chris Shannon (1994) "Monotone comparative statics," *Econometrica: Journal of the Econometric Society*, 157–180.

Mirrlees, James A (1971) "An exploration in the theory of optimum income taxation," *The review of economic studies*, 38 (2), 175–208.

Ramey, Garey and Valerie A Ramey (2009) "The rug rat race,"Technical report, National Bureau of Economic Research.

Restuccia, Diego and Carlos Urrutia (2004) "Intergenerational persistence of earnings: The role of early and college education," *American Economic Review*, 94 (5), 1354–1378.

Stantcheva, Stefanie (2015) "Optimal income, education, and bequest taxes in an intergenerational model,"Technical report, National Bureau of Economic Research.

——— (2017) "Optimal taxation and human capital policies over the life cycle," *Journal of Political Economy*, 125 (6), 1931–1990.

——— (2020) "Dynamic taxation," *Annual Review of Economics*, 12 (1), 801–831.

A Solution of the relaxed recursive problem: Hamiltonian approach

In this appendix, we derive the first-order conditions for the relaxed recursive planner's problem using optimal-control techniques in the spirit of Mirrlees (1971), following Koeniger and Prat (2018). Throughout, we maintain the regularity, interiority, and first-order-approach assumptions stated in the main text and in Appendix F. In particular, we assume that the relevant policy functions are sufficiently smooth, that differentiation under the integral sign is valid, and that $u(\cdot)$ is strictly increasing so that u^{-1} is well defined on the relevant range.

Throughout, we use the reduced-form representation of period utility obtained after substituting out labor supply. Specifically, we write

$$U(c, y, h, \theta, h') \equiv u(c) - v(y, h, \theta) + \varphi(h'),$$

where $v(y, h, \theta)$ denotes the reduced-form disutility associated with producing output y, given current human capital h and ability θ. Accordingly,

$$U_c(c, y, h, \theta, h') = u_c(c), \qquad U_y(c, y, h, \theta, h') = -v_y(y, h, \theta), \qquad U_{h'}(c, y, h, \theta, h') = \varphi_{h'}(h').$$

Using (11), we substitute consumption in the planner's objective by the implicit function

$$c(\omega(\theta) - \beta V'(\theta),\, y(\theta),\, h,\, \theta,\, h'(\theta))\,.$$

We treat $\omega(\theta)$ as a state variable, with law of motion given by the recursive envelope condition (12), and associated costate $\mu(\theta)$. We associate multiplier λ with the promise-keeping constraint (13) and multiplier γ with the marginal-promise-keeping constraint (14). As in Farhi and Werning (2013), the envelope conditions for the value function are

$$\Gamma_V(V, \Phi, h, \theta_-, t) = \lambda, \qquad \Gamma_\Phi(V, \Phi, h, \theta_-, t) = \gamma. \tag{48}$$

To simplify notation, we use a prime $'$ for next-period objects and an underscore $_-$ for the previous-period type. In particular, for a reported current type θ,

$$V'(\theta) \equiv \int_\Theta \omega(\theta')\, dF(\theta' \mid \theta), \qquad \Phi'(\theta) \equiv \int_\Theta \omega(\theta') \frac{\partial f(\theta' \mid \theta)}{\partial \theta}\, d\theta'.$$

Hamiltonian. The current-value Hamiltonian associated with the type-indexed control problem is

$$\begin{aligned}\mathcal{H} = &\Big[c(\omega(\theta) - \beta V'(\theta),\, y(\theta),\, h,\, \theta,\, h'(\theta)) + g(h'(\theta), h) - y(\theta) + q\,\Gamma(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)\Big] f(\theta \mid \theta_-)\\ &+ \lambda\Big[V - \omega(\theta) f(\theta \mid \theta_-)\Big] + \gamma\Big[\Phi - \omega(\theta)\frac{\partial f(\theta \mid \theta_-)}{\partial \theta_-}\Big]\\ &+ \mu(\theta)\left[\frac{\partial U(c(\omega(\theta) - \beta V'(\theta),\, y(\theta),\, h,\, \theta,\, h'(\theta))\,,\, y(\theta),\, h,\, \theta,\, h'(\theta))}{\partial \theta} + \beta \Phi'(\theta)\right].\end{aligned} \tag{49}$$

The first line collects current resource costs and discounted continuation costs. The second line imposes promise keeping and marginal promise keeping. The third line imposes the recursive envelope condition through the costate $\mu(\theta)$. Because consumption has been substituted out using the promise-keeping relation, the Hamiltonian is written in terms of current resource costs rather than direct utility from consumption.

First-order conditions. The first-order conditions with respect to the controls $\big(V'(\theta), \Phi'(\theta), h'(\theta), y(\theta)\big)$ are necessary conditions for an interior optimum of the relaxed problem. Differentiating the Hamiltonian with respect to each control and dividing by $f(\theta \mid \theta_-) > 0$, we obtain the following conditions. For notational convenience, $U(\cdot)$ denotes

$$U(c(\omega(\theta) - \beta V'(\theta),\, y(\theta),\, h,\, \theta,\, h'(\theta))\,,\, y(\theta),\, h,\, \theta,\, h'(\theta))\,.$$

First, with respect to $V'(\theta)$,

$$q\frac{\partial \Gamma(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)}{\partial V'(\theta)} = -\frac{\partial c(\theta)}{\partial V'(\theta)} - \frac{\mu(\theta)}{f(\theta \mid \theta_-)}\frac{\partial^2 U(\cdot)}{\partial \theta\, \partial c(\theta)}\frac{\partial c(\theta)}{\partial V'(\theta)}. \tag{50}$$

Second, with respect to $\Phi'(\theta)$,

$$q\frac{\partial \Gamma(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)}{\partial \Phi'(\theta)} = -\frac{\beta \mu(\theta)}{f(\theta \mid \theta_-)}. \tag{51}$$

Third, with respect to $h'(\theta)$,

$$\begin{aligned} q\frac{\partial \Gamma(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)}{\partial h'(\theta)} &= -\frac{\partial c(\theta)}{\partial h'(\theta)} - \frac{\partial g(h'(\theta), h)}{\partial h'(\theta)} \\ &\quad - \frac{\mu(\theta)}{f(\theta \mid \theta_-)}\left(\frac{\partial^2 U(\cdot)}{\partial \theta\, \partial c(\theta)}\frac{\partial c(\theta)}{\partial h'(\theta)} + \frac{\partial^2 U(\cdot)}{\partial \theta\, \partial h'(\theta)}\right). \end{aligned} \tag{52}$$

Fourth, with respect to $y(\theta)$,

$$0 = \frac{\partial c(\theta)}{\partial y(\theta)} - 1 + \frac{\mu(\theta)}{f(\theta \mid \theta_-)}\left[\frac{\partial^2 U(\cdot)}{\partial \theta\, \partial c(\theta)}\frac{\partial c(\theta)}{\partial y(\theta)} + \frac{\partial^2 U(\cdot)}{\partial \theta\, \partial y(\theta)}\right]. \tag{53}$$

Equation (53) is the planner's intratemporal condition for output.

Costate dynamics and boundary conditions. The costate $\mu(\theta)$ satisfies the adjoint equation associated with the state variable $\omega(\theta)$. Differentiating the Hamiltonian with respect to $\omega(\theta)$ yields

$$\frac{\partial \mu(\theta)}{\partial \theta} = -\left[\frac{\partial c(\theta)}{\partial \omega(\theta)} - \lambda - \gamma\frac{\frac{\partial f(\theta \mid \theta_-)}{\partial \theta_-}}{f(\theta \mid \theta_-)} + \frac{\mu(\theta)}{f(\theta \mid \theta_-)}\frac{\partial^2 U(\cdot)}{\partial \theta\, \partial c(\theta)}\frac{\partial c(\theta)}{\partial \omega(\theta)}\right] f(\theta \mid \theta_-). \tag{54}$$

The corresponding natural boundary conditions are

$$\mu(\underline{\theta}) = 0, \qquad \mu(\bar{\theta}) = 0. \tag{55}$$

Closed-form derivatives of $c(\theta)$. Using the promise-keeping relation (11), consumption can be written as

$$c(\omega(\theta) - \beta V'(\theta),\, y(\theta),\, h,\, \theta,\, h'(\theta)) = u^{-1}(\omega(\theta) - \beta V'(\theta) + v(y(\theta), h, \theta) - \varphi(h'(\theta)))\,. \tag{56}$$

Applying the inverse function theorem yields

$$\frac{\partial c(\theta)}{\partial \omega(\theta)} = \frac{1}{u_c(c(\theta))}, \tag{57}$$

$$\frac{\partial c(\theta)}{\partial V'(\theta)} = -\frac{\beta}{u_c(c(\theta))}, \tag{58}$$

$$\frac{\partial c(\theta)}{\partial y(\theta)} = \frac{v_y(y(\theta), h, \theta)}{u_c(c(\theta))}, \tag{59}$$

$$\frac{\partial c(\theta)}{\partial h} = \frac{v_h(y(\theta), h, \theta)}{u_c(c(\theta))}, \tag{60}$$

$$\frac{\partial c(\theta)}{\partial h'(\theta)} = -\frac{\varphi_{h'}(h'(\theta))}{u_c(c(\theta))}. \tag{61}$$

Condition for $V'(\theta)$. Since ability enters period utility only through the reduced-form labor disutility term $v(y, h, \theta)$, utility from consumption is type-independent. Hence,

$$\frac{\partial^2 U(\cdot)}{\partial \theta \, \partial c(\theta)} = 0.$$

Using this property in (50), we obtain

$$q \frac{\partial \Gamma(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)}{\partial V'(\theta)} = -\frac{\partial c(\theta)}{\partial V'(\theta)} = \frac{\beta}{u_c(c(\theta))}. \tag{62}$$

Now define

$$\lambda'(\theta) \equiv \Gamma_V(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)\,.$$

Using (48), equation (62) implies

$$\frac{1}{u_c(c(\theta))} = \frac{q}{\beta} \lambda'(\theta). \tag{63}$$

Equation (63) links current reciprocal marginal utility to the continuation-value multiplier. Evaluating the costate law of motion at the upper bound and using the boundary condition $\mu(\bar{\theta}) = 0$ yields the recursive restriction on λ, from which the reciprocal Euler equation follows:

$$\frac{1}{u_c(c_t)} = \frac{q}{\beta} \mathbb{E}_t \left[\frac{1}{u_c(c_{t+1})} \right]. \tag{64}$$

Condition for $y(\theta)$. Under $U_{\theta c}(\cdot) = 0$, the output first-order condition (53) simplifies to

$$1 - \frac{\partial c(\theta)}{\partial y(\theta)} = \frac{\mu(\theta)}{f(\theta \mid \theta_-)} \frac{\partial^2 U(\cdot)}{\partial \theta \, \partial y(\theta)}. \tag{65}$$

Because

$$U(c, y, h, \theta, h') = u(c) - v(y, h, \theta) + \varphi(h'),$$

we have

$$\frac{\partial^2 U(\cdot)}{\partial \theta \, \partial y(\theta)} = -\frac{\partial^2 v(y(\theta), h, \theta)}{\partial \theta \, \partial y(\theta)}. \tag{66}$$

Substituting (66) into (65), and using (59), gives

$$-\frac{\mu(\theta)}{f(\theta \mid \theta_-)}\frac{\partial^2 v(y(\theta),h,\theta)}{\partial\theta\,\partial y(\theta)} = 1 - \frac{v_y(y(\theta),h,\theta)}{u_c(c(\theta))}. \tag{67}$$

Equation (67) is the planner's intratemporal optimality condition for output.

Envelope condition with respect to current human capital h. To derive the condition for $h'(\theta)$, we first compute the envelope condition with respect to inherited human capital h. Differentiating the Bellman equation for $\Gamma(V,\Phi,h,\theta_-,t)$ with respect to h, while holding optimal controls fixed, yields

$$\begin{aligned}\Gamma_h(V,\Phi,h,\theta_-,t) &= \int_\Theta \left(\frac{\partial c(\theta)}{\partial h} + \frac{\partial g(h'(\theta),h)}{\partial h}\right) dF(\theta \mid \theta_-) \\ &\quad + \int_\Theta \mu(\theta)\left(\frac{\partial^2 U(\cdot)}{\partial\theta\,\partial c}\frac{\partial c(\theta)}{\partial h} + \frac{\partial^2 U(\cdot)}{\partial\theta\,\partial h}\right) d\theta.\end{aligned} \tag{68}$$

Since

$$\frac{\partial^2 U(\cdot)}{\partial\theta\,\partial c} = 0, \qquad \frac{\partial^2 U(\cdot)}{\partial\theta\,\partial h} = -\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial h},$$

and using (60), we obtain

$$\begin{aligned}\Gamma_h(V,\Phi,h,\theta_-,t) &= \int_\Theta \left(\frac{v_h(y(\theta),h,\theta)}{u_c(c(\theta))} + \frac{\partial g(h'(\theta),h)}{\partial h}\right) dF(\theta \mid \theta_-) \\ &\quad - \int_\Theta \mu(\theta)\frac{\partial^2 v(y(\theta),h,\theta)}{\partial\theta\,\partial h}\, d\theta.\end{aligned} \tag{69}$$

Condition for $h'(\theta)$. We now return to (52). Since ability enters utility neither through consumption nor through the altruistic component $\varphi(h')$, we have

$$\frac{\partial^2 U(\cdot)}{\partial\theta\,\partial c} = 0, \qquad \frac{\partial^2 U(\cdot)}{\partial\theta\,\partial h'} = 0.$$

Therefore, (52) simplifies to

$$q\,\frac{\partial\Gamma(V'(\theta),\Phi'(\theta),h'(\theta),\theta,t+1)}{\partial h'(\theta)} = -\frac{\partial c(\theta)}{\partial h'(\theta)} - \frac{\partial g(h'(\theta),h)}{\partial h'(\theta)}. \tag{70}$$

Substituting (61) into (70) yields

$$q\,\frac{\partial\Gamma(V'(\theta),\Phi'(\theta),h'(\theta),\theta,t+1)}{\partial h'(\theta)} = \frac{\varphi_{h'}(h'(\theta))}{u_c(c(\theta))} - \frac{\partial g(h'(\theta),h)}{\partial h'(\theta)}. \tag{71}$$

Applying (69) at period $t+1$, with inherited human capital equal to $h'(\theta)$ and previous-period type equal to the current report θ, yields

$$\begin{aligned}\frac{\partial\Gamma(V'(\theta),\Phi'(\theta),h'(\theta),\theta,t+1)}{\partial h'(\theta)} &= \int_\Theta \left(\frac{v_h(y',h',\theta')}{u_c(c'(\theta'))} + \frac{\partial g(h''(\theta'),h')}{\partial h'}\right) dF(\theta' \mid \theta) \\ &\quad - \int_\Theta \mu'(\theta')\frac{\partial^2 v(y',h',\theta')}{\partial\theta'\,\partial h'}\, d\theta'.\end{aligned} \tag{72}$$

Substituting (72) into (71) yields

$$q\left(\int_{\Theta}\left(\frac{v_h(y',h',\theta')}{u_c(c'(\theta'))}+\frac{\partial g(h''(\theta'),h')}{\partial h'}\right)dF(\theta'\mid\theta)-\int_{\Theta}\mu'(\theta')\frac{\partial^2 v(y',h',\theta')}{\partial\theta'\,\partial h'}\,d\theta'\right)=\frac{\varphi_{h'}(h'(\theta))}{u_c(c(\theta))}-\frac{\partial g(h'(\theta),h)}{\partial h'(\theta)}. \tag{73}$$

Equation (73) is the planner's optimality condition for human capital investment in the relaxed problem.

Simplified costate dynamics. Finally, under

$$\frac{\partial^2 U(\cdot)}{\partial\theta\,\partial c}=0,$$

the costate dynamics (54) simplify to

$$\frac{\partial\mu(\theta)}{\partial\theta}=-\left[\frac{\partial c(\theta)}{\partial\omega(\theta)}-\lambda-\gamma\frac{\frac{\partial f(\theta\mid\theta_-)}{\partial\theta_-}}{f(\theta\mid\theta_-)}\right]f(\theta\mid\theta_-). \tag{74}$$

Using (57), we obtain

$$\frac{\partial\mu(\theta)}{\partial\theta}=-\left[\frac{1}{u_c(c(\theta))}-\lambda-\gamma\frac{\frac{\partial f(\theta\mid\theta_-)}{\partial\theta_-}}{f(\theta\mid\theta_-)}\right]f(\theta\mid\theta_-). \tag{75}$$

Integrating (75) from $\underline{\theta}$ to θ, and using $\mu(\underline{\theta})=0$, yields

$$\mu(\theta)=\int_{\underline{\theta}}^{\theta}\left[-\frac{1}{u_c(c(x))}+\lambda+\gamma\frac{\frac{\partial f(x\mid\theta_-)}{\partial\theta_-}}{f(x\mid\theta_-)}\right]dF(x\mid\theta_-). \tag{76}$$

Summary. The Hamiltonian approach yields three key necessary conditions for the relaxed planner's allocation.

First, equations (63) and (64) link reciprocal marginal utility to the continuation-value multiplier and imply the reciprocal Euler equation.

Second, equation (67) gives the planner's intratemporal optimality condition for output.

Third, equation (73) gives the planner's optimality condition for human capital investment. This condition is the key input for the derivation of the constrained-efficient human-capital wedge reported in the main text.

B Laissez-faire optimality conditions

This appendix derives the laissez-faire first-order conditions for bequests, labor supply, and human capital investment. Throughout, the dynasty's per-period utility function is

$$U(c_t,l_t,h_{t+1})=u(c_t)-v(l_t)+\varphi(h_{t+1}).$$

Accordingly,

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} = \frac{\partial u(c_t)}{\partial c_t}, \qquad \frac{\partial U(c_t, l_t, h_{t+1})}{\partial l_t} = -\frac{\partial v(l_t)}{\partial l_t}, \qquad \frac{\partial U(c_t, l_t, h_{t+1})}{\partial h_{t+1}} = \frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}}.$$

At date t, the dynasty enters the period with financial assets b_t, human capital h_t, and realized privately observed ability θ_t. Here, b_t denotes financial assets carried into period t, which can be interpreted as bequests received from the previous generation. Conditional on these state variables, the dynasty chooses current consumption c_t, labor supply l_t, next-generation human capital h_{t+1}, and next-period assets b_{t+1}.

The only remaining uncertainty concerns next period's ability θ_{t+1}. Under the Markov assumption for ability, the conditional distribution of θ_{t+1} depends only on the current realization θ_t. Accordingly, conditioning on the full history is equivalent, for one-step-ahead expectations, to conditioning on θ_t alone. Throughout this appendix, we therefore use the shorthand

$$\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot \mid \theta_t],$$

where expectations are taken with respect to the conditional distribution $dF(\theta_{t+1} \mid \theta_t)$.

To simplify notation, we suppress the dependence of policy functions on θ_t whenever no confusion arises. Because choices are made after observing θ_t, the first-order conditions derived below hold pointwise in θ_t.

B.1 Dynasty problem

The dynasty solves the Bellman problem

$$W_t(b_t, h_t, \theta_t) = \max_{c_t, l_t, h_{t+1}, b_{t+1}} \left\{ U(c_t, l_t, h_{t+1}) + \beta \mathbb{E}_t[W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})] \right\}, \tag{77}$$

subject to

$$c_t + b_{t+1} + g(h_{t+1}, h_t) = (1+r)b_t + y_t, \tag{78}$$

$$y_t = Y(h_t, \theta_t, l_t). \tag{79}$$

There are no taxes, subsidies, or borrowing constraints in the laissez-faire economy.

B.2 Lagrangian

Let λ_t denote the multiplier on the within-period budget constraint (78). The associated Lagrangian is

$$\begin{aligned} \mathcal{L}_t = {} & U(c_t, l_t, h_{t+1}) + \beta \mathbb{E}_t[W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})] \\ & + \lambda_t \Big[(1+r)b_t + Y(h_t, \theta_t, l_t) - c_t - b_{t+1} - g(h_{t+1}, h_t) \Big]. \end{aligned} \tag{80}$$

B.3 First-order conditions

We restrict attention to interior allocations and derive the corresponding first-order conditions for the dynasty's choice variables.

Consumption. Differentiating (80) with respect to c_t gives

$$\frac{\partial \mathcal{L}_t}{\partial c_t} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} - \lambda_t = 0. \tag{81}$$

Therefore,

$$\lambda_t = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}. \tag{82}$$

Next-period assets. Differentiating (80) with respect to b_{t+1} yields

$$\frac{\partial \mathcal{L}_t}{\partial b_{t+1}} = -\lambda_t + \beta \mathbb{E}_t \left[\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial b_{t+1}} \right] = 0. \tag{83}$$

Substituting (82) into (83) gives

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} = \beta \mathbb{E}_t \left[\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial b_{t+1}} \right]. \tag{84}$$

Human capital. Differentiating (80) with respect to h_{t+1} yields

$$\frac{\partial \mathcal{L}_t}{\partial h_{t+1}} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial h_{t+1}} - \lambda_t \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} + \beta \mathbb{E}_t \left[\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial h_{t+1}} \right] = 0. \tag{85}$$

Using (82), we obtain

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial h_{t+1}} + \beta \mathbb{E}_t \left[\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial h_{t+1}} \right]. \tag{86}$$

Labor supply. Differentiating (80) with respect to l_t gives

$$\frac{\partial \mathcal{L}_t}{\partial l_t} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial l_t} + \lambda_t \frac{\partial Y(h_t, \theta_t, l_t)}{\partial l_t} = 0. \tag{87}$$

Substituting (82) into (87) yields

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} \frac{\partial Y(h_t, \theta_t, l_t)}{\partial l_t} + \frac{\partial U(c_t, l_t, h_{t+1})}{\partial l_t} = 0. \tag{88}$$

B.4 Envelope conditions

We next derive the envelope conditions that allow us to eliminate derivatives of the continuation value function from the first-order conditions.

B.4.1 Envelope condition with respect to financial assets

Differentiating the Bellman equation (77) with respect to current financial assets b_t and applying the envelope theorem implies that all terms involving derivatives of the optimal policy functions vanish. The effect of b_t on the value function therefore operates only through its impact on current resources via the budget constraint. Holding optimal choices fixed, differentiating (78) yields

$$\frac{\partial W_t(b_t, h_t, \theta_t)}{\partial b_t} = \lambda_t \frac{\partial}{\partial b_t} \left[(1+r) b_t \right] = \lambda_t (1+r). \tag{89}$$

By (82), this implies

$$\frac{\partial W_t(b_t, h_t, \theta_t)}{\partial b_t} = (1+r)\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}. \tag{90}$$

Applying the same argument one period ahead gives

$$\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial b_{t+1}} = (1+r)\frac{\partial U(c_{t+1}, l_{t+1}, h_{t+2})}{\partial c_{t+1}}. \tag{91}$$

Substituting (91) into (84) yields

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t} = \beta(1+r)\,\mathbb{E}_t\left[\frac{\partial U(c_{t+1}, l_{t+1}, h_{t+2})}{\partial c_{t+1}}\right]. \tag{92}$$

Using separability, this becomes

$$\frac{\partial u(c_t)}{\partial c_t} = \beta(1+r)\,\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]. \tag{93}$$

B.4.2 Envelope condition with respect to human capital

Differentiating the Bellman equation (77) with respect to h_t and applying the envelope theorem implies that all terms involving derivatives of the optimal policy functions vanish. The effect of h_t on the value function therefore operates only through its impact on current resources via output and education costs.

Holding optimal choices fixed, this implies that differentiating (78) and using (79) yields

$$\frac{\partial c_t}{\partial h_t} = \frac{\partial Y(h_t, \theta_t, l_t)}{\partial h_t} - \frac{\partial g(h_{t+1}, h_t)}{\partial h_t}. \tag{94}$$

Since current-period utility does not depend directly on h_t in our specification, it follows that

$$\frac{\partial W_t(b_t, h_t, \theta_t)}{\partial h_t} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}\frac{\partial c_t}{\partial h_t} \tag{95}$$

$$= \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}\left(\frac{\partial Y(h_t, \theta_t, l_t)}{\partial h_t} - \frac{\partial g(h_{t+1}, h_t)}{\partial h_t}\right). \tag{96}$$

Applying the same argument one period ahead, with state variable h_{t+1}, gives

$$\frac{\partial W_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})}{\partial h_{t+1}} = \frac{\partial U(c_{t+1}, l_{t+1}, h_{t+2})}{\partial c_{t+1}}\left(\frac{\partial Y(h_{t+1}, \theta_{t+1}, l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}}\right). \tag{97}$$

Substituting (97) into (86) yields

$$\frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}\frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} = \frac{\partial U(c_t, l_t, h_{t+1})}{\partial h_{t+1}} + \beta\mathbb{E}_t\left[\frac{\partial U(c_{t+1}, l_{t+1}, h_{t+2})}{\partial c_{t+1}}\left(\frac{\partial Y(h_{t+1}, \theta_{t+1}, l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}}\right)\right]. \tag{98}$$

Using separability, this becomes

$$\frac{\partial u(c_t)}{\partial c_t}\frac{\partial g(h_{t+1},h_t)}{\partial h_{t+1}} = \frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}} + \beta \mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\left(\frac{\partial Y(h_{t+1},\theta_{t+1},l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2},h_{t+1})}{\partial h_{t+1}}\right)\right]. \tag{99}$$

For the terminal period, the continuation value vanishes and the expectation term drops out. Accordingly, the optimality condition for human capital reduces to a static trade-off between marginal utility from investment and its marginal cost.

B.4.3 Labor supply condition

Finally, using separability, the labor-supply first-order condition (88) can be rewritten as

$$\frac{\partial u(c_t)}{\partial c_t}\frac{\partial Y(h_t,\theta_t,l_t)}{\partial l_t} = \frac{\partial v(l_t)}{\partial l_t}. \tag{100}$$

B.5 Summary

Collecting the previous results, the laissez-faire allocation satisfies

$$\frac{\partial u(c_t)}{\partial c_t} = \beta(1+r)\,\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right], \tag{101}$$

$$\frac{\partial u(c_t)}{\partial c_t}\frac{\partial g(h_{t+1},h_t)}{\partial h_{t+1}} = \frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}} + \beta \mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\left(\frac{\partial Y(h_{t+1},\theta_{t+1},l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2},h_{t+1})}{\partial h_{t+1}}\right)\right], \tag{102}$$

$$\frac{\partial u(c_t)}{\partial c_t}\frac{\partial Y(h_t,\theta_t,l_t)}{\partial l_t} = \frac{\partial v(l_t)}{\partial l_t}. \tag{103}$$

These are the laissez-faire conditions used in the main text to define the bequest, labor, and human-capital wedges.

C Proofs

Proof of Remark 1: Reciprocal Euler equation. The reciprocal Euler equation follows from the planner's first-order condition with respect to promised utility in the relaxed recursive problem. Consider the first-order condition for $V'(\theta_t)$ given in (50).

Since ability θ_t enters preferences only through the disutility of labor, utility from consumption is independent of θ_t. It follows that

$$\frac{\partial^2 U(c_t,l_t,h_{t+1})}{\partial \theta_t\,\partial c_t} = 0,$$

so that all cross-derivative terms involving c_t and θ_t vanish. Under this assumption, (50) simplifies to

$$q\,\Gamma_{V,t+1}(\theta_t) = -\frac{\partial c_t(\theta_t)}{\partial V'(\theta_t)}. \tag{104}$$

Next, use the envelope condition for promised utility, which implies

$$\Gamma_{V,t}(\theta_t) = \lambda_t(\theta_t).$$

Evaluating this condition one period ahead yields

$$\Gamma_{V,t+1}(\theta_{t+1}) = \lambda_{t+1}(\theta_{t+1}).$$

Substituting into (104) gives

$$q\,\lambda_{t+1}(\theta_{t+1}) = -\frac{\partial c_t(\theta_t)}{\partial V'(\theta_t)}. \tag{105}$$

To evaluate the derivative on the right-hand side, recall that consumption is defined implicitly by (56). Differentiating that expression with respect to $V'(\theta_t)$ and applying the implicit function theorem yields

$$\frac{\partial c_t(\theta_t)}{\partial V'(\theta_t)} = -\frac{\beta}{\frac{\partial u(c_t)}{\partial c_t}}. \tag{106}$$

Substituting (106) into (105) gives

$$\frac{1}{\frac{\partial u(c_t)}{\partial c_t}} = \frac{q}{\beta}\,\lambda_{t+1}(\theta_{t+1}). \tag{107}$$

Using the envelope condition for consumption at $t+1$, we have

$$\lambda_{t+1}(\theta_{t+1}) = \frac{1}{\frac{\partial u(c_{t+1})}{\partial c_{t+1}}}.$$

Substituting into (107) and taking expectations conditional on θ_t yields the reciprocal Euler equation:

$$\frac{1}{\frac{\partial u(c_t)}{\partial c_t}} = \frac{q}{\beta}\mathbb{E}_t\left[\frac{1}{\frac{\partial u(c_{t+1})}{\partial c_{t+1}}}\right]. \tag{108}$$

This condition coincides with Atkeson and Lucas (1992) and Koeniger and Prat (2018) in the absence of human capital externalities.

Proof of Proposition 1: Optimal human capital condition in the relaxed problem. The proposition follows from the optimality conditions of the planner's relaxed problem derived in Section A. We provide the key steps linking these conditions to equation (16).

Step 1: Human capital first-order condition. The planner's first-order condition with respect to $h'(\theta)$, given in (70), can be written as

$$q\,\Gamma_h(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1) = \frac{\varphi_{h'}(h'(\theta))}{u_c(c(\theta))} - \frac{\partial g(h'(\theta), h)}{\partial h'(\theta)}, \tag{109}$$

where we have used $U_{\theta c} = 0$, $U_{\theta h'} = 0$, and substituted (61). Here $\Gamma_h(\cdot, t+1)$ denotes the partial derivative of the planner's cost function with respect to inherited human capital, evaluated one period ahead.

Step 2: Substitution of the envelope condition. Evaluating the envelope condition (69) at period $t+1$, with inherited human capital equal to $h'(\theta)$ and previous-period type equal to θ, yields an explicit expression for $\Gamma_h(\cdot, t+1)$ in terms of marginal utilities, marginal products, and the costate variable $\mu_{t+1}(\theta_{t+1})$. Substituting this expression into (109) delivers equation (16).

Step 3: Representation of the costate variable. The integral representation of the costate variable $\mu_{t+1}(\theta_{t+1})$ in (17) follows from integrating the costate dynamics (54) from $\underline{\theta}$ to θ, using the lower boundary condition $\mu(\underline{\theta}) = 0$ to pin down the constant of integration.

Step 4: Boundary conditions. The boundary conditions

$$\mu_{t+1}(\underline{\theta}) = 0, \qquad \mu_{t+1}(\overline{\theta}) = 0$$

are the boundary conditions associated with the type-indexed control problem. The lower condition pins down the constant of integration in Step 3; the upper condition pins down the multiplier λ.

Step 5: Determination of λ. Evaluating (17) at $\theta = \overline{\theta}$ and imposing $\mu_{t+1}(\overline{\theta}) = 0$ yields

$$\lambda = \mathbb{E}_t\left[\left(\frac{\partial u(c_t(\theta_t))}{\partial c_t}\right)^{-1}\right],$$

which completes the proof.

C.0.1 Definition of wedges as implicit marginal taxes and subsidies

Throughout this appendix, let $v(y, h, \theta)$ denote the reduced-form disutility obtained by substituting the production technology $y = Y(h, \theta, l)$ into $v(l)$. By definition,

$$v_y(y, h, \theta) \equiv \frac{\partial v(y, h, \theta)}{\partial y} = \frac{v_l(l)}{Y_l(h, \theta, l)},$$

evaluated at the labor input l that generates output y.

Wedges are defined as *local marginal distortions* applied to the dynasty's laissez-faire optimality conditions. For each margin $j \in \{b, l, h\}$, the wedge $\tau_j(\theta^t)$ is chosen so that the distorted laissez-faire condition coincides with the planner's marginal condition evaluated at the same allocation. These planner conditions correspond to the first-order conditions of the relaxed planner's problem described in the main text.

Bequest wedge $\tau_b(\theta^t)$. Introduce a local marginal tax on bequests such that the dynasty internalizes only a fraction $(1-\tau_b(\theta^t))$ of the gross return $(1+r)$ on transfers to the next period. The distorted Euler equation is

$$\frac{\partial u(c_t)}{\partial c_t} = \beta(1+r)\,(1-\tau_b(\theta^t))\,\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]. \tag{110}$$

Solving (110) for $\tau_b(\theta^t)$ and using $q = (1+r)^{-1}$ yields

$$\tau_b(\theta^t) \equiv 1 - \frac{q}{\beta}\,\frac{\frac{\partial u(c_t)}{\partial c_t}}{\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]}. \tag{111}$$

Labor wedge $\tau_l(\theta^t)$. In the laissez-faire economy, the intratemporal optimality condition can be written in reduced form as

$$\frac{\partial u(c_t)}{\partial c_t} = \frac{\partial v(y_t, h_{t-1}, \theta_t)}{\partial y_t}.$$

Introduce a local marginal tax on labor income such that the dynasty internalizes only a fraction (1 −

$\tau_l(\theta^t)$) of the marginal return to labor. The distorted labor condition becomes

$$\frac{\partial v(y_t, h_{t-1}, \theta_t)}{\partial y_t} = (1 - \tau_l(\theta^t)) \frac{\partial u(c_t)}{\partial c_t}. \tag{112}$$

Rearranging (112) yields

$$\tau_l(\theta^t) \equiv 1 - \frac{\frac{\partial v(y_t, h_{t-1}, \theta_t)}{\partial y_t}}{\frac{\partial u(c_t)}{\partial c_t}}. \tag{113}$$

Hence, the labor wedge measures the proportional gap between the marginal disutility of output (in reduced form) and the marginal utility of consumption.

Human capital wedge $\tau_h(\theta^t)$. For periods $t < T$, the laissez-faire optimality condition for human capital investment is

$$\frac{\partial u(c_t)}{\partial c_t} \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} = \frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}} + \beta \, \mathbb{E}_t \left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}} \left(\frac{\partial Y(h_{t+1}, \theta_{t+1}, l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}} \right) \right].$$

We model the human capital wedge as a local distortion to the marginal cost of investment. Suppose the dynasty faces the effective marginal cost

$$(1 + \tau_h(\theta^t)) \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}}.$$

The distorted condition is

$$\frac{\partial u(c_t)}{\partial c_t} (1 + \tau_h(\theta^t)) \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} - \frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}} = \beta \, \mathbb{E}_t \left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}} \left(\frac{\partial Y(h_{t+1}, \theta_{t+1}, l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}} \right) \right]. \tag{114}$$

Solving (114) for $\tau_h(\theta^t)$ yields

$$\begin{aligned} \tau_h(\theta^t) \equiv{}& \frac{\beta}{\frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}}} \mathbb{E}_t \left[\frac{\frac{\partial u(c_{t+1})}{\partial c_{t+1}}}{\frac{\partial u(c_t)}{\partial c_t}} \left(\frac{\partial Y(h_{t+1}, \theta_{t+1}, l_{t+1})}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}} \right) \right] \\ &+ \frac{\frac{\partial \varphi(h_{t+1})}{\partial h_{t+1}}}{\frac{\partial u(c_t)}{\partial c_t} \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}}} - 1. \end{aligned} \tag{115}$$

A positive value of $\tau_h(\theta^t)$ corresponds to an implicit tax on human capital investment, increasing the effective marginal cost, while a negative value corresponds to a subsidy.

C.0.2 Constrained-efficient wedges

The *constrained-efficient wedges* are obtained by evaluating the definitions (111)–(115) at the allocation chosen by the planner under private information.

Constrained-efficient bequest wedge $\tau_b^*(\theta^t)$. At the planner allocation, the optimality conditions imply the reciprocal Euler equation

$$\frac{1}{\frac{\partial u(c_t)}{\partial c_t}} = \frac{q}{\beta}\,\mathbb{E}_t\left[\frac{1}{\frac{\partial u(c_{t+1})}{\partial c_{t+1}}}\right]. \tag{116}$$

Evaluating (111) at the planner allocation gives

$$\tau_b^*(\theta^t) = 1 - \frac{q}{\beta}\frac{\frac{\partial u(c_t)}{\partial c_t}}{\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]}. \tag{117}$$

Using (116) to eliminate the current marginal utility yields

$$\tau_b^*(\theta^t) = 1 - \frac{1}{\mathbb{E}_t\left[\frac{1}{\frac{\partial u(c_{t+1})}{\partial c_{t+1}}}\right]\,\mathbb{E}_t\left[\frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]}. \tag{118}$$

Since $\frac{\partial u(c_{t+1})}{\partial c_{t+1}} > 0$, Jensen's inequality implies

$$\mathbb{E}_t[X]\,\mathbb{E}_t\left[\frac{1}{X}\right] \geq 1 \qquad \text{for } X > 0,$$

hence $\tau_b^*(\theta^t) \geq 0$, with strict inequality whenever consumption varies across states.

Constrained-efficient labor wedge $\tau_l^*(\theta^t)$. Evaluating the definition (113) at the planner allocation yields

$$\tau_l^*(\theta^t) = 1 - \frac{\frac{\partial v(y_t,h_t,\theta_t)}{\partial y_t}}{\frac{\partial u(c_t)}{\partial c_t}}. \tag{119}$$

The planner's first-order condition for output implies

$$1 = \frac{\frac{\partial v(y_t,h_t,\theta_t)}{\partial y_t}}{\frac{\partial u(c_t)}{\partial c_t}} + \frac{\mu(\theta^t)}{f(\theta_t \mid \theta_{t-1})}\frac{\partial^2 v(y_t,h_t,\theta_t)}{\partial \theta_t\,\partial y_t}.$$

Rearranging gives

$$\frac{\frac{\partial v(y_t,h_t,\theta_t)}{\partial y_t}}{\frac{\partial u(c_t)}{\partial c_t}} = 1 - \frac{\mu(\theta^t)}{f(\theta_t \mid \theta_{t-1})}\frac{\partial^2 v(y_t,h_t,\theta_t)}{\partial \theta_t\,\partial y_t}.$$

Substituting into (119) yields the constrained-efficient labor wedge

$$\tau_l^*(\theta^t) = -\frac{\mu(\theta^t)}{f(\theta_t \mid \theta_{t-1})}\frac{\partial^2 v(y_t,h_t,\theta_t)}{\partial \theta_t\,\partial y_t}. \tag{120}$$

This expression corresponds to the standard Mirrleesian formula expressed in reduced form.

Constrained-efficient human capital wedge $\tau_h^*(\theta^t)$.

Proof of Proposition 4. The proof proceeds in three steps. First, we express the human capital wedge as the difference between the planner's optimality condition and the distorted laissez-faire Euler equation. Second, we simplify the intertemporal-transfer component Δ_b using the bequest wedge. Third, we combine the labor-mediated term with the raw incentive term and recover Δ_i.

Step 1: A raw decomposition of τ_h.

By Definition 1, the human capital wedge τ_h is defined by the condition that the distorted laissez-faire Euler equation for human capital coincides with the planner's allocation:

$$\frac{\partial u(c)}{\partial c}(1+\tau_h)\frac{\partial g(h',h)}{\partial h'} - \frac{\partial \varphi(h')}{\partial h'} = \beta\,\mathbb{E}\left[\frac{\partial u(c')}{\partial c'}\left(\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right)\right]. \tag{121}$$

Equation (121) is algebraically equivalent to Definition 115 after multiplying both sides of the definition of τ_h by $\frac{\partial u(c)}{\partial c}\frac{\partial g(h',h)}{\partial h'}$.

Solving (121) for τ_h yields

$$\tau_h = \frac{1}{\frac{\partial g(h',h)}{\partial h'}}\left[\frac{\frac{\partial \varphi(h')}{\partial h'}}{\frac{\partial u(c)}{\partial c}} + \frac{\beta}{\frac{\partial u(c)}{\partial c}}\mathbb{E}\left[\frac{\partial u(c')}{\partial c'}\left(\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right)\right]\right] - 1. \tag{122}$$

Next, recall the planner's first-order condition for $h'(\theta)$ from Proposition **??**:

$$\frac{\partial g(h',h)}{\partial h'} = \frac{\frac{\partial \varphi(h')}{\partial h'}}{\frac{\partial u(c)}{\partial c}} - q\,\mathbb{E}\left[\frac{\frac{\partial v(y',h',\theta')}{\partial h'}}{\frac{\partial u(c')}{\partial c'}} + \frac{\partial g(h'',h')}{\partial h'}\right] + q\,\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial \theta'\,\partial h'}\right]. \tag{123}$$

Substituting (123) into (122), the term $\frac{\partial \varphi(h')/\partial h'}{\partial u(c)/\partial c}$ cancels. Moreover, using

$$\frac{\partial v(y,h,\theta)}{\partial h} = -\frac{\partial v(y,h,\theta)}{\partial y}\frac{\partial y}{\partial h}, \tag{124}$$

which follows from $y = A(\theta,h)l$ and $v(y,h,\theta) = \boldsymbol{v}(y/A(\theta,h))$, we obtain

$$\begin{aligned}\tau_h = {} & \underbrace{\frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\frac{\partial y'}{\partial h'}\left(1 - \frac{\frac{\partial v(y',h',\theta')}{\partial y'}}{\frac{\partial u(c')}{\partial c'}}\right)\right]}_{\equiv \Delta_l} \\ & + \underbrace{\frac{1}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\left(\beta\frac{\frac{\partial u(c')}{\partial c'}}{\frac{\partial u(c)}{\partial c}} - q\right)\left(\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right)\right]}_{\equiv \Delta_b} \\ & - \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial \theta'\,\partial h'}\right].\end{aligned} \tag{125}$$

The term Δ_l captures the labor-mediated component of the wedge. The final expectation in (125) is the raw incentive effect of human capital on future implementability.

Step 2: Simplifying Δ_b.

Write Δ_b as

$$\Delta_b = \frac{1}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\left(\beta\frac{\frac{\partial u(c')}{\partial c'}}{\frac{\partial u(c)}{\partial c}} - q\right)\left(\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right)\right]. \tag{126}$$

Adding and subtracting the product of expectations gives

$$\Delta_b = \frac{1}{\frac{\partial g(h',h)}{\partial h'}} \mathbb{E}\left[\frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right] \mathbb{E}\left[\beta \frac{\frac{\partial u(c')}{\partial c'}}{\frac{\partial u(c)}{\partial c}} - q\right] + \frac{\beta}{\frac{\partial g(h',h)}{\partial h'} \frac{\partial u(c)}{\partial c}} \operatorname{Cov}\left(\frac{\partial u(c')}{\partial c'}, \frac{\partial y'}{\partial h'} - \frac{\partial g(h'',h')}{\partial h'}\right). \tag{127}$$

From Definition 1,

$$\tau_b^* = 1 - \frac{q}{\beta} \frac{\frac{\partial u\left(c_t(\theta^t)\right)}{\partial c_t}}{\mathbb{E}_t\left[\frac{\partial u\left(c_{t+1}(\theta^{t+1})\right)}{\partial c_{t+1}}\right]},$$

which rearranges to

$$\frac{\beta\, \mathbb{E}_t\left[\frac{\partial u\left(c_{t+1}(\theta^{t+1})\right)}{\partial c_{t+1}}\right]}{\frac{\partial u(c_t(\theta^t))}{\partial c_t}} = \frac{q}{1-\tau_b^*}.$$

Since $\frac{\partial u(c_t(\theta^t))}{\partial c_t}$ is measurable with respect to period-t information, it factors out of $\mathbb{E}_t[\,\cdot\,]$, so subtracting q from both sides gives

$$\mathbb{E}\left[\beta \frac{\frac{\partial u\left(c_{t+1}(\theta^{t+1})\right)}{\partial c_{t+1}}}{\frac{\partial u(c_t(\theta^t))}{\partial c_t}} - q\right] = q \frac{\tau_b^*}{1-\tau_b^*}.$$

Substituting into the first term of (127) yields exactly (30).

Step 3: Recovering Δ_i.

We now combine Δ_l with the final term in (125). By Lemma 2, for $v(y,h,\theta) = \boldsymbol{v}(y/A(\theta,h))$,

$$\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial h} = -\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial y}\frac{\partial y}{\partial h} + \frac{d\boldsymbol{v}(l)}{dl}\frac{l}{A(\theta,h)^2}\frac{\partial A(\theta,h)}{\partial\theta}\frac{\partial A(\theta,h)}{\partial h}\Xi(\theta,h). \tag{128}$$

Applying (128) next period and taking expectations gives

$$\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial\theta'\,\partial h'}\right] = -\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial\theta'\,\partial y'}\frac{\partial y'}{\partial h'}\right] + \mathbb{E}\left[\Xi(\theta',h')\frac{d\boldsymbol{v}(l'(\theta'))}{dl'}\frac{l'(\theta')}{A(\theta',h')^2}\frac{\partial A(\theta',h')}{\partial\theta'}\frac{\partial A(\theta',h')}{\partial h'}\frac{\mu'(\theta')}{f(\theta'|\theta)}\right]. \tag{129}$$

From the planner's first-order condition for output $y(\theta)$ (Proposition 3),

$$-\frac{\mu(\theta)}{f(\theta \mid \theta_{t-1})}\frac{\partial^2 v(y,h,\theta)}{\partial\theta\,\partial y} = 1 - \frac{\frac{\partial v(y,h,\theta)}{\partial y}}{\frac{\partial u(c)}{\partial c}}. \tag{130}$$

Multiplying both sides of (130) by $(\partial y'/\partial h')f(\theta'|\theta)$ and integrating over Θ gives

$$\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial\theta'\,\partial y'}\frac{\partial y'}{\partial h'}\right] = \mathbb{E}\left[\left(\frac{\frac{\partial v(y',h',\theta')}{\partial y'}}{\frac{\partial u(c')}{\partial c'}} - 1\right)\frac{\partial y'}{\partial h'}\right]. \tag{131}$$

Substituting (129) and (131) into (125), the sum of Δ_l and the final term in (125) becomes

$$\begin{aligned}
&\Delta_l - \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial \theta'\,\partial h'}\right] \\
&\quad= \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\frac{\partial y'}{\partial h'}\left(1 - \frac{\frac{\partial v(y',h',\theta')}{\partial y'}}{\frac{\partial u(c')}{\partial c'}}\right)\right] \\
&\qquad+ \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\left(\frac{\frac{\partial v(y',h',\theta')}{\partial y'}}{\frac{\partial u(c')}{\partial c'}} - 1\right)\frac{\partial y'}{\partial h'}\right] \\
&\qquad- \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\Xi(\theta',h')\,\frac{d\boldsymbol{v}(l'(\theta'))}{dl'}\,\frac{l'(\theta')}{A(\theta',h')^2}\,\frac{\partial A(\theta',h')}{\partial \theta'}\,\frac{\partial A(\theta',h')}{\partial h'}\,\frac{\mu'(\theta')}{f(\theta'|\theta)}\right].
\end{aligned} \tag{132}$$

The first two expectations cancel identically. Hence,

$$\Delta_l - \frac{q}{\frac{\partial g(h',h)}{\partial h'}}\mathbb{E}\left[\frac{\mu'(\theta')}{f(\theta'|\theta)}\frac{\partial^2 v(y',h',\theta')}{\partial \theta'\,\partial h'}\right] = \Delta_i. \tag{133}$$

Substituting this identity back into (125) establishes

$$\tau_h^* = \Delta_b + \Delta_i.$$

This completes the proof. □

Proof of Lemma 2. Since $l = y/A(\theta,h)$, we have

$$v(y,h,\theta) = \boldsymbol{v}(l).$$

The relevant derivatives are

$$\frac{\partial v}{\partial y} = \frac{d\boldsymbol{v}(l)}{dl}\,\frac{1}{A(\theta,h)}, \tag{134}$$

$$\frac{\partial v}{\partial h} = -\frac{d\boldsymbol{v}(l)}{dl}\,\frac{y}{A(\theta,h)^2}\,A_h(\theta,h), \tag{135}$$

$$\frac{\partial v}{\partial \theta} = -\frac{d\boldsymbol{v}(l)}{dl}\,\frac{y}{A(\theta,h)^2}\,A_\theta(\theta,h). \tag{136}$$

Differentiate (135) with respect to θ. Using

$$\frac{\partial l}{\partial \theta} = -\frac{l\,A_\theta(\theta,h)}{A(\theta,h)},$$

we obtain

$$
\begin{aligned}
\frac{\partial^2 v}{\partial\theta\,\partial h} &= -\frac{\partial}{\partial\theta}\left[\frac{d\boldsymbol{v}(l)}{dl}\,\frac{y}{A(\theta,h)^2}\,A_h(\theta,h)\right] \\
&= \frac{d^2\boldsymbol{v}(l)}{dl^2}\,\frac{l^2\,A_\theta(\theta,h)\,A_h(\theta,h)}{A(\theta,h)^2} - \frac{d\boldsymbol{v}(l)}{dl}\,\frac{l}{A(\theta,h)}\,A_{\theta h}(\theta,h) \\
&\quad + 2\,\frac{d\boldsymbol{v}(l)}{dl}\,\frac{l\,A_\theta(\theta,h)\,A_h(\theta,h)}{A(\theta,h)^2}.
\end{aligned}
\tag{137}
$$

Next, differentiating (134) with respect to θ, we get

$$
\frac{\partial^2 v}{\partial\theta\,\partial y} = -\frac{d^2\boldsymbol{v}(l)}{dl^2}\,\frac{l\,A_\theta(\theta,h)}{A(\theta,h)^2} - \frac{d\boldsymbol{v}(l)}{dl}\,\frac{A_\theta(\theta,h)}{A(\theta,h)^2}.
$$

Multiplying by $\partial y/\partial h = l\,A_h(\theta,h)$, it follows that

$$
-\frac{\partial^2 v}{\partial\theta\,\partial y}\,\frac{\partial y}{\partial h} = \frac{d^2\boldsymbol{v}(l)}{dl^2}\,\frac{l^2\,A_\theta(\theta,h)\,A_h(\theta,h)}{A(\theta,h)^2} + \frac{d\boldsymbol{v}(l)}{dl}\,\frac{l\,A_\theta(\theta,h)\,A_h(\theta,h)}{A(\theta,h)^2}.
\tag{138}
$$

Subtracting (138) from (137) yields

$$
\frac{\partial^2 v}{\partial\theta\,\partial h} + \frac{\partial^2 v}{\partial\theta\,\partial y}\,\frac{\partial y}{\partial h} = \frac{d\boldsymbol{v}(l)}{dl}\,\frac{l}{A(\theta,h)^2}\,A_\theta(\theta,h)\,A_h(\theta,h)\left[1 - \frac{A_{\theta h}(\theta,h)\,A(\theta,h)}{A_\theta(\theta,h)\,A_h(\theta,h)}\right].
\tag{139}
$$

Rearranging gives (34). □

C.0.3 Summary

Equations (111)–(115) define policy wedges as local distortions of the laissez-faire optimality conditions. Evaluating these at the planner allocation yields the constrained-efficient wedges: (118) for the bequest wedge, (120) for the labor wedge, and (??) together with (??) and (??) for the decomposition of the human capital wedge.

D Numerical solution of the family problem

This appendix describes the numerical solution of the decentralized dynasty problem. We solve the household problem using the endogenous gridpoint method (EGM) proposed by Carroll (2006) and extended by Hintermaier and Koeniger (2010). The method is based on the first-order conditions and is well suited for stochastic dynamic problems with continuous state and control variables, including occasionally binding inequality constraints.

D.1 Household problem

For periods $t < T$, a household with state (b_t, h_t, θ_t) chooses next-period assets b_{t+1}, next-period human capital h_{t+1}, and labor supply l_t. Current consumption c_t is then determined from the budget constraint. The value function satisfies

$$
\widehat{W}_t(b_t,h_t,\theta_t) = \max_{b_{t+1},h_{t+1},l_t}\left\{U_t(c_t,l_t,h_{t+1}) + \beta\int_\Theta \widehat{W}_{t+1}(b_{t+1},h_{t+1},\theta_{t+1})\,dF(\theta_{t+1}\mid\theta_t)\right\},
\tag{140}
$$

subject to

$$c_t = (1+r)b_t - T^b(b_t) + y_t - T^y(y_t) - b_{t+1} - g(h_{t+1}, h_t), \tag{141}$$

where

$$y_t \equiv Y(h_t, \theta_t, l_t), \tag{142}$$

the borrowing constraint

$$b_{t+1} \geq \max\{-\phi\, g(h_{t+1}, h_t),\ \underline{b}\}, \tag{143}$$

and the lower bound on next-period human capital

$$h_{t+1} \geq 1. \tag{144}$$

Here $g(h_{t+1}, h_t)$ denotes the cost of investing in next-generation human capital h_{t+1}, given parental human capital h_t. The borrowing limit in (143) allows the household to finance a fraction ϕ of education expenditures through debt, subject to an exogenous floor $\underline{b}$.

D.2 Lagrangian and first-order conditions

Let $\nu_t \geq 0$ denote the Kuhn–Tucker multiplier on (143) and $\eta_t \geq 0$ the multiplier on (144). The complementary-slackness conditions are

$$\nu_t \geq 0, \quad b_{t+1} - \max\{-\phi\, g(h_{t+1}, h_t),\ \underline{b}\} \geq 0, \quad \nu_t\Big(b_{t+1} - \max\{-\phi\, g(h_{t+1}, h_t),\ \underline{b}\}\Big) = 0, \tag{145}$$

$$\eta_t \geq 0, \quad h_{t+1} - 1 \geq 0, \quad \eta_t(h_{t+1} - 1) = 0. \tag{146}$$

The Lagrangian is

$$\begin{aligned}\mathcal{L}_t(b_{t+1}, h_{t+1}, l_t; \nu_t, \eta_t) = U_t(c_t, l_t, h_{t+1}) + \beta \int_\Theta \widehat{W}_{t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t) \\ + \nu_t\Big(b_{t+1} - \max\{-\phi\, g(h_{t+1}, h_t),\ \underline{b}\}\Big) + \eta_t(h_{t+1} - 1),\end{aligned} \tag{147}$$

where c_t is given by (141).

Because the borrowing constraint involves the kinked expression $\max\{-\phi\, g(h_{t+1}, h_t),\ \underline{b}\}$, the first-order condition with respect to h_{t+1} is branch-specific.

First-order condition with respect to b_{t+1}. Within each branch, the constraint is differentiable and the first-order condition is

$$U_{c,t}(c_t, l_t, h_{t+1}) = \beta \int_\Theta \widehat{W}_{b,t+1}(b_{t+1}, h_{t+1}, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t) + \nu_t. \tag{148}$$

First-order condition with respect to h_{t+1}. Using $\partial c_t / \partial h_{t+1} = -g_{h'}(h_{t+1}, h_t)$ from (141), where

$$g_{h'}(h_{t+1}, h_t) \equiv \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}},$$

we obtain

$$-U_{c,t}(c_t,l_t,h_{t+1})\,g_{h'}(h_{t+1},h_t)+U_{h,t}(c_t,l_t,h_{t+1})+\beta\int_\Theta \widehat{W}_{h,t+1}(b_{t+1},h_{t+1},\theta_{t+1})\,dF(\theta_{t+1}\mid\theta_t)+\eta_t=0, \tag{149}$$

when the active borrowing branch is $b_{t+1}=\underline{b}$. By contrast, when the active branch is $b_{t+1}=-\phi\,g(h_{t+1},h_t)$, differentiation of the constraint contributes the additional term $\nu_t\phi g_{h'}(h_{t+1},h_t)$, so that

$$-U_{c,t}(c_t,l_t,h_{t+1})\,g_{h'}(h_{t+1},h_t)+U_{h,t}(c_t,l_t,h_{t+1})+\beta\int_\Theta \widehat{W}_{h,t+1}(b_{t+1},h_{t+1},\theta_{t+1})\,dF(\theta_{t+1}\mid\theta_t)+\nu_t\,\phi\,g_{h'}(h_{t+1},h_t)+\eta_t=0. \tag{150}$$

First-order condition with respect to l_t. Using

$$\frac{\partial c_t}{\partial l_t}=Y_l(h_t,\theta_t,l_t)\big[1-T^{y\prime}(y_t)\big],$$

where

$$Y_l(h_t,\theta_t,l_t)\equiv\frac{\partial Y(h_t,\theta_t,l_t)}{\partial l_t},\qquad T^{y\prime}(y_t)\equiv\frac{\partial T^y(y_t)}{\partial y_t},$$

we obtain

$$U_{c,t}(c_t,l_t,h_{t+1})\,Y_l(h_t,\theta_t,l_t)\big[1-T^{y\prime}(y_t)\big]+U_{l,t}(c_t,l_t,h_{t+1})=0. \tag{151}$$

D.3 Auxiliary object Λ_t

Substituting (148) into the first-order condition with respect to h_{t+1} yields the auxiliary object

$$\begin{aligned}\Lambda_t(b_{t+1},h_{t+1},h_t,\theta_t)\equiv U_{h,t}(c_t,l_t,h_{t+1})+\beta\int_\Theta\Big[&\widehat{W}_{h,t+1}(b_{t+1},h_{t+1},\theta_{t+1})\\ &-\widehat{W}_{b,t+1}(b_{t+1},h_{t+1},\theta_{t+1})\,g_{h'}(h_{t+1},h_t)\Big]\,dF(\theta_{t+1}\mid\theta_t).\end{aligned} \tag{152}$$

In the branch $b_{t+1}=-\phi\,g(h_{t+1},h_t)$, (148) and (150) imply

$$\Lambda_t(b_{t+1},h_{t+1},h_t,\theta_t)=\nu_t\,g_{h'}(h_{t+1},h_t)\big(1-\phi\big)-\eta_t, \tag{153}$$

whereas in the branch $b_{t+1}=\underline{b}$ we have

$$\Lambda_t(b_{t+1},h_{t+1},h_t,\theta_t)=\nu_t\,g_{h'}(h_{t+1},h_t)-\eta_t. \tag{154}$$

The function Λ_t summarizes the net marginal value of human capital investment beyond what is captured by the bequest margin.

D.4 Terminal period

In the terminal period T, there is no continuation value:

$$\widehat{W}_{T+1}=0.$$

By assumption, households cannot die in debt, so $b_{T+1} \geq 0$. Since utility from bequests is absent, the optimum satisfies

$$b_{T+1} = 0. \tag{155}$$

In the terminal period, there is no subsequent generation within the planning horizon. Households therefore choose only current consumption c_T and labor supply l_T and do not invest in children's human capital. The terminal-period utility is

$$U_T(c_T, l_T) = \ln c_T - \frac{l_T^\alpha}{\alpha}. \tag{156}$$

The budget constraint is

$$c_T = (1+r)b_T - T^b(b_T) + y_T - T^y(y_T), \tag{157}$$

where

$$y_T \equiv Y(h_T, \theta_T, l_T). \tag{158}$$

Let λ_T denote the multiplier on (157). The first-order conditions are

$$\frac{1}{c_T} = \lambda_T, \tag{159}$$

$$l_T^{\alpha-1} = \lambda_T\, Y_l(h_T, \theta_T, l_T)\big[1 - T^{y\prime}(y_T)\big]. \tag{160}$$

Under the parametric specification

$$y_T = l_T \theta_T^\xi h_T^{1-\xi}, \qquad T^y(y_T) = y_T - \delta y_T^{1-t_y},$$

so that

$$T^{y\prime}(y_T) = 1 - \delta(1-t_y) y_T^{-t_y},$$

equation (160) becomes

$$l_T^{\alpha-1+t_y} = \lambda_T(1-t_y)\delta(\theta_T^\xi h_T^{1-\xi})^{1-t_y}. \tag{161}$$

Substituting (161) into (157) yields a single nonlinear equation in c_T, which we solve by root-finding. Given c_T, optimal labor supply l_T is recovered from (161).

D.5 Algorithm for periods $t < T$

For $t < T$, we follow Koeniger and Prat (2018). Fix grids

$$G_{b_{t+1}} = \{b_{t+1}^1, \ldots, b_{t+1}^I\}, \qquad G_{h_t} = \{h_t^1, \ldots, h_t^J\}, \qquad G_{\theta_t} = \{\theta_t^1, \ldots, \theta_t^K\}.$$

Throughout the numerical implementation, we use the regularized cost function

$$g(h_{t+1}, h_t) = \tilde{\kappa}\big(\ln(h_{t+1} + h^*)\big)^{\varsigma_1}\big(\ln(h_t + h^*)\big)^{\varsigma_2}, \quad h^* = 0.001, \tag{162}$$

which ensures $g_{h'}(h_{t+1}, h_t) \neq 0$ throughout the state space. The shift h^* is a numerical regularization that does not affect the economic logic of the problem.

For each $(b_{t+1}^i, h_t^j, \theta_t^k)$, the feasible range of h_{t+1} depends on the active borrowing branch. When $b_{t+1}^i \leq 0$, the lower bound on h_{t+1} is obtained by inverting the branch $b^i_{t+1} = -\phi\, g(h_{t+1}, h^j_t)$:

$$\underline{h}^{ij}_{t+1} = g^{-1}\left(-\frac{b^i_{t+1}}{\phi},\, h^j_t\right). \tag{163}$$

The corresponding upper bound $\overline{h}^{ijk}_{t+1}$ is determined, when needed, by

$$\Lambda_t(b^i_{t+1}, \overline{h}^{ijk}_{t+1}, h^j_t, \theta^k_t) = 0, \qquad \overline{h}^{ijk}_{t+1} \leq h^J_t. \tag{164}$$

D.6 Cases with occasionally binding constraints

We distinguish four cases based on the binding status of (143) and (144).

Case 1: $b^i_{t+1} > 0$ ($\nu_t = 0$).

[(a)]

1. If $\Lambda_t(b^i_{t+1}, 1, h^j_t, \theta^k_t) > 0$, then $h_{t+1} > 1$ and we find the root of $\Lambda_t = 0$.
2. If $\Lambda_t(b^i_{t+1}, 1, h^j_t, \theta^k_t) \leq 0$, then $h_{t+1} = 1$ and $\eta_t = -\Lambda_t(b^i_{t+1}, 1, h^j_t, \theta^k_t) \geq 0$.
3. If no root exists in $(1, h^J_t]$, we extend the upper bound of the grid until one is found.

Case 2: $b^i_{t+1} < 0$, active branch $b_{t+1} = -\phi\, g(h_{t+1}, h_t)$.

[(a)]

1. Compute $\underline{h}^{ij}_{t+1}$ from (163).
2. If $\Lambda_t(b^i_{t+1}, \underline{h}^{ij}_{t+1}, h^j_t, \theta^k_t) \geq 0$, the endogenous borrowing limit is binding at the optimum: $h_{t+1} = \underline{h}^{ij}_{t+1}$ and $\nu_t > 0$.
3. If $\Lambda_t(b^i_{t+1}, \underline{h}^{ij}_{t+1}, h^j_t, \theta^k_t) < 0$, the solution lies above the lower bound implied by the borrowing branch, and we find the root of $\Lambda_t = 0$ for $h_{t+1} > \underline{h}^{ij}_{t+1}$.

Case 3: $b^i_{t+1} > 0$ and $h_{t+1} = 1$. Then $\nu_t = 0$ and

$$\eta_t = -\Lambda_t(b^i_{t+1}, 1, h^j_t, \theta^k_t) > 0.$$

Case 4: both constraints bind. This case corresponds to the corner at which the human-capital constraint binds, $h_{t+1} = 1$, and the borrowing choice lies at the intersection of the relevant lower bounds implied by (143). Denote this value by $b_{t+1,\text{corner}}$.

Current consumption is then given by the budget constraint:

$$c_{ijk}(l_t) = (1+r)b^i_t - T^b(b^i_t) + Y(h^j_t, \theta^k_t, l_t) - T^y\big(Y(h^j_t, \theta^k_t, l_t)\big) - b_{t+1,\text{corner}} - g(1, h^j_t). \tag{165}$$

Substituting (165) into (151) yields optimal labor supply. The multipliers ν_t and η_t are then recovered from (148) and the appropriate branch-specific first-order condition for h_{t+1}.

D.7 Endogenous gridpoints and interpolation

For Cases 1–3, consumption is recovered by inverting (148). Since the utility function is separable in consumption and labor, the marginal utility of consumption does not depend on labor supply, so inversion yields c_t without conditioning on l_t:

$$u_c(c_t^{nijk}) = \beta \int_\Theta \widehat{W}_{b,t+1}(b_{t+1}^i, h_{t+1}^n, \theta_{t+1})\, dF(\theta_{t+1} \mid \theta_t^k) + \nu_t^{nijk}. \tag{166}$$

Labor supply is then recovered from (151):

$$-U_{l,t}(c_t^{nijk}, l_t^{nijk}, h_{t+1}^n) = u_c(c_t^{nijk})\, Y_l(h_t^j, \theta_t^k, l_t^{nijk})\big[1 - T^{y\prime}(y_t^{nijk})\big]. \tag{167}$$

The endogenous gridpoint for current assets is defined implicitly by

$$b_t^{nijk} = \frac{b_{t+1}^i + T^b(b_t^{nijk}) + c_t^{nijk} + g(h_{t+1}^n, h_t^j) - y_t^{nijk} + T^y(y_t^{nijk})}{1+r}, \tag{168}$$

where

$$y_t^{nijk} \equiv Y(h_t^j, \theta_t^k, l_t^{nijk}).$$

When T^b is linear in b_t, (168) has a closed-form solution; otherwise it is solved pointwise by root-finding.

This step defines the endogenous grid for current assets. After constructing the endogenous grid, we interpolate the policy functions $c_t(b_t, h_t, \theta_t)$, $l_t(b_t, h_t, \theta_t)$, and $h_{t+1}(b_t, h_t, \theta_t)$ onto the exogenous grid $G_{b_t} = G_{b_{t+1}}$. The model is then solved by backward induction. The implementation is carried out in Matlab.

E Numerical solution of the social planner problem

This appendix describes the numerical algorithm used to solve the social planner's problem. The implementation is in Matlab. The computational strategy follows Farhi and Werning (2013) and Koeniger and Prat (2018). We therefore provide a concise description of the numerical procedure and refer the reader to those papers for additional implementation details.

We solve the planner's problem using the first-order approach and backward induction. The state variables of the recursive problem are $(V_t, \Phi_t, h_t, \theta_{t-1}, t)$, as in (10). The terminal condition is that there is no continuation beyond period T. Hence,

$$V_{T+1}(\theta^T) = 0, \qquad \Phi_{T+1}(\theta^T) = 0. \tag{169}$$

Accordingly, the numerical solution separates into two cases: the terminal period T and the recursive step for periods $t < T$.

E.1 Terminal period T

In the terminal period, continuation values vanish and the planner solves a static allocation problem conditional on the inherited state (h_T, θ_{T-1}) and on the promised objects inherited from period $T-1$.

Using (11) together with (169), the continuation utility simplifies to

$$\omega(\theta^T) = U\big(c_T(\theta^T), y_T(\theta^T), h_T(\theta^{T-1}), \theta_T, h_{T+1}(\theta^T)\big)\,. \tag{170}$$

Output is determined from the planner's intratemporal condition (67). Specializing to period T,

$$1 - \frac{v_y(y(\theta), h(\theta_-), \theta)}{u_c(c(\theta))} + \frac{\mu(\theta)}{f(\theta \mid \theta_-)} v_{\theta y}(y, h, \theta) = 0. \tag{171}$$

For the logarithmic utility specification,

$$u_c(c(\theta)) = \frac{1}{c(\theta)}. \tag{172}$$

Using (11), consumption can be recovered from $\omega(\theta^T)$ as

$$\ln c_T(\theta^T) = \omega(\theta^T) + \frac{\left(\frac{y(\theta^T)}{A(h^T,\theta^T)}\right)^{\alpha}}{\alpha} - \vartheta \ln h^{T+1}, \tag{173}$$

$$c_T(\theta^T) = \exp\left(\omega(\theta^T) + \frac{\left(\frac{y(\theta^T)}{A(h^T,\theta^T)}\right)^{\alpha}}{\alpha} - \vartheta \ln h^{T+1}\right). \tag{174}$$

Substituting into (171) yields a nonlinear equation in output, which is solved numerically at each grid point.

Given terminal allocations, the costate $\mu(\theta)$ is updated using (75), while $\omega(\theta)$ is updated from the envelope condition

$$\frac{\partial \omega(\theta)}{\partial \theta} = -\frac{\partial v(y(\theta), h, \theta)}{\partial \theta} + \beta \int_{\Theta} \omega(\theta') \frac{\partial f(\theta' \mid \theta)}{\partial \theta} d\theta'. \tag{175}$$

Finally, we compute the return to human capital using the planner's condition (73). These returns serve as inputs for the backward recursion.

E.2 Recursive step for periods $t < T$

For each period $t < T$, the planner's problem is solved conditional on the state $(V_t, \Phi_t, h_t, \theta_{t-1}, t)$.

Step 1: Output. Output is computed from the planner's intratemporal condition

$$1 - \frac{v_y(y(\theta), h(\theta_-), \theta)}{u_c(c(\theta))} + \frac{\mu(\theta)}{f(\theta \mid \theta_-)} v_{\theta y}(y, h, \theta) = 0. \tag{176}$$

This equation is solved numerically for $y(\theta)$, taking as given $\mu(\theta)$ and continuation objects.

Step 2: Consumption. Consumption is pinned down by the reciprocal Euler equation

$$\frac{1}{u_c(c(\theta))} = \frac{q}{\beta} \lambda'(\theta), \tag{177}$$

where $\lambda'(\theta) = \Gamma_V(V'(\theta), \Phi'(\theta), h'(\theta), \theta, t+1)$. In the numerical implementation, $\lambda'(\theta)$ is obtained by interpolation over the continuation-value grid.

Step 3: Human capital. Optimal human capital investment is computed from the planner's condition (73). Numerically, we solve for $h'(\theta)$ as the root of

$$-q\left(\int_{\Theta}\left(\frac{v_h(y'(\theta'),h',\theta')}{u_c(c'(\theta'))}+g_h(h''(\theta'),h')\right)dF(\theta'\mid\theta)-\int_{\Theta}\mu'(\theta')v_{\theta'h'}(y',h',\theta')\,d\theta'\right)$$
$$+\frac{\varphi'(h'(\theta))}{u_c(c(\theta))}-g_{h'}(h'(\theta),h)=0. \tag{178}$$

Crucially, the first term is integrated with respect to the transition probability $dF(\theta'\mid\theta)$, while the incentive term involving the costate $\mu'(\theta')$ is integrated with respect to the Lebesgue measure $d\theta'$, consistently with the planner's first-order conditions.

Step 4: Updating allocations. Given $y(\theta)$, $h'(\theta)$, and continuation values, consumption follows from (11). Under logarithmic utility,

$$c_t(\theta^t)=\exp\left(\omega(\theta^t)-\beta V_{t+1}(\theta^t)+\frac{\left(\frac{y(\theta^t)}{A(h^t,\theta^t)}\right)^{\alpha}}{\alpha}-\vartheta\ln h^{t+1}\right). \tag{179}$$

We then update $\omega(\theta)$, $\mu(\theta)$, and continuation objects (V_{t+1},Φ_{t+1}), and proceed backward to period $t-1$.

At each iteration, we also compute the implied returns to human capital and the corresponding wedges defined in Appendix C.

F Validity of the first-order approach

This appendix provides the formal argument underlying Lemma 1.

F.1 Step 1: one-shot deviations are sufficient

By Lemma 1 of Kapička (2013), it is without loss of generality to restrict attention to strategies that deviate from truthful reporting in a single period t and report truthfully in all other periods. For any history θ^{t-1} and report $\sigma\in\Theta$, define the one-shot deviation utility:

$$\omega_t^r(\theta^{t-1},\theta_t;\sigma)=U\big(c_t(\theta^{t-1},\sigma),\,y_t(\theta^{t-1},\sigma),\,h_t(\theta^{t-1}),\,\theta_t,\,h_{t+1}(\theta^{t-1},\sigma)\big)$$
$$+\,\beta\int_{\Theta}\omega_{t+1}(\theta^{t-1},\sigma,\theta_{t+1})\,dF(\theta_{t+1}|\theta_t). \tag{180}$$

Global incentive compatibility then requires $\omega_t(\theta^t)\geq\omega_t^r(\theta^{t-1},\theta_t;\sigma)$ for all $\sigma\in\Theta$ and all $\theta^t\in\Theta^t$.

F.2 Step 2: the envelope condition

By Assumption 1(i), ω_t^r is absolutely continuous in θ_t. Since truthful reporting is optimal, the first-order condition $\partial\omega_t^r/\partial\sigma|_{\sigma=\theta_t}=0$ holds at an interior solution, and the envelope theorem of Milgrom and Segal (2002) (Theorem 2) applies. Differentiating $\omega_t(\theta^t)=\omega_t^r(\theta^{t-1},\theta_t;\theta_t)$ with respect to θ_t yields the recursive envelope condition (9):

$$\frac{\partial \omega_t(\theta^t)}{\partial \theta_t} = \frac{\partial U(\cdot)}{\partial \theta_t} + \beta \int_\Theta \omega_{t+1}(\theta^t, \theta_{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)}{\partial \theta_t} \, d\theta_{t+1}. \tag{181}$$

F.3 Step 3: single-crossing is preserved under altruism

The key step is to verify that the introduction of parental altruism $\varphi(h_{t+1})$ does not violate the single-crossing condition that drives global incentive compatibility. Computing the cross-partial derivative of the deviation utility (180) with respect to the report σ and the true type θ_t:

$$\frac{\partial^2 \omega_t^r}{\partial \sigma \, \partial \theta_t} = -\frac{\partial^2 v\big(l_t(\theta^{t-1}, \sigma), h_t, \theta_t\big)}{\partial \sigma \, \partial \theta_t} + \underbrace{\frac{\partial^2 \varphi\big(h_{t+1}(\theta^{t-1}, \sigma)\big)}{\partial \sigma \, \partial \theta_t}}_{=0} + \underbrace{\beta \int_\Theta \omega_{t+1} \frac{\partial^2 f(\theta_{t+1}|\theta_t)}{\partial \sigma \, \partial \theta_t} \, d\theta_{t+1}}_{=0}. \tag{182}$$

The second term vanishes because $\varphi(h_{t+1}(\theta^{t-1}, \sigma))$ depends on the report σ but not on the true type θ_t: parental altruism enters utility through the *chosen* human capital level, which is a function of the report, not of the true type. The third term vanishes because $f(\theta_{t+1}|\theta_t)$ depends on θ_t but not on σ. Therefore (182) reduces to:

$$\frac{\partial^2 \omega_t^r}{\partial \sigma \, \partial \theta_t} = -\frac{\partial^2 v\big(l_t(\theta^{t-1}, \sigma), h_t, \theta_t\big)}{\partial \sigma \, \partial \theta_t} \;\geq\; 0, \tag{183}$$

where the inequality follows from Assumption 1(ii). Hence parental altruism $\varphi(h_{t+1})$ leaves the single-crossing condition intact.

F.4 Step 4: local implies global incentive compatibility

Condition (183) implies that $\omega_t^r(\theta^{t-1}, \theta_t; \sigma)$ satisfies increasing differences in (σ, θ_t). By Assumption 1(iii), the allocation is monotone, so the optimal report is non-decreasing in θ_t. Together, these two properties imply that the local first-order condition $\partial \omega_t^r / \partial \sigma|_{\sigma=\theta_t} = 0$ characterizes a global maximum (see Milgrom and Shannon (1994), Theorem 4):

$$\omega_t^r(\theta^{t-1}, \theta_t; \theta_t) \;\geq\; \omega_t^r(\theta^{t-1}, \theta_t; \sigma) \quad \forall \, \sigma \in \Theta.$$

The allocation is therefore globally incentive compatible, which completes the proof of Lemma 1. □

G Implementation of the Constrained-Efficient Allocation: Derivations

This section studies how the constrained-efficient allocation can be implemented in a decentralized economy using simple policy instruments. The central question is whether there exists a (potentially complex) tax system such that, when dynasties face these taxes, their optimal choices replicate the allocation chosen by the social planner.

Because the problem is dynamic and features multiple choice margins, it is convenient to interpret wedges as implicit taxes and subsidies (Stantcheva, 2020). Importantly, the mapping between wedges and explicit tax schedules is not one-to-one. As emphasized by Stantcheva (2020), wedges characterize distortions along a single dimension, holding other choices fixed at the optimum. When agents choose multiple actions, joint deviations may arise, complicating the implementation problem. For a comprehensive overview of recent advances in the implementation of optimal tax systems, see Stantcheva (2020).

We focus on marginal taxes and show how they can be chosen to replicate the wedges derived in the planner's problem. The decentralized economy is characterized by a general history-dependent tax system

$$T_t(b_t, y_t, h_{t+1}, \theta^{t-1}),$$

which conditions on observable outcomes and the reported history of types. Following Albanesi and Sleet (2006), Farhi and Werning (2013), and Koeniger and Prat (2018), incentive compatibility is enforced through off-path punishments.

Given this tax system, each dynasty solves

$$\max_{\{c_s(\theta^s), l_s(\theta^s), h_{s+1}(\theta^s), b_{s+1}(\theta^s)\}} \mathbb{E}_t \left[\sum_{s=t}^{T} \beta^{s-t} U(c_s, l_s, h_{s+1}) \right], \tag{184}$$

subject to the budget constraint

$$b_{t+1} + c_t + g(h_{t+1}, h_t) = (1+r)b_t + y_t - T_t(b_t, y_t, h_{t+1}, \theta^{t-1}), \tag{185}$$

and the production technology

$$y_t = Y(h_t, \theta_t, l_t).$$

Labor income taxation. The first-order condition with respect to labor implies that the marginal labor income tax equals the labor wedge. Specifically, combining the labor FOC with the tax schedule yields

$$\frac{\partial T_t}{\partial y_t} = \tau_{l,t}, \tag{186}$$

which mirrors the result in Koeniger and Prat (2018). Thus, the optimal labor income tax directly implements the labor wedge.

Bequest taxation. Similarly, the Euler equation for bequests implies that the marginal tax on bequests is pinned down by the bequest wedge. Using the envelope condition,

$$\frac{\partial W(b_t, h_t)}{\partial b_t} = \left(1 + r - \frac{\partial T_t}{\partial b_t}\right) \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t},$$

and the first-order condition for b_{t+1}, we obtain

$$\frac{\partial U(c_t)}{\partial c_t} = \beta \mathbb{E}_t \left[\left(1 + r - \frac{\partial T_{t+1}}{\partial b_{t+1}}\right) \frac{\partial U(c_{t+1})}{\partial c_{t+1}} \right]. \tag{187}$$

As in Koeniger and Prat (2018), the bequest tax must satisfy this condition ex post for each reported type. Otherwise, dynasties would have incentives to deviate by increasing bequests and reducing labor supply in the next generation.

Education taxation. Finally, the marginal tax (or subsidy) on children's human capital investment is determined by the human capital wedge $\tau_{h,t}$. Using the first-order condition for h_{t+1} and the expression for the wedge derived in Section 4, the optimal education policy is chosen so that the decentralized Euler equation for human capital coincides with its constrained-efficient counterpart.

Taken together, these results show that the constrained-efficient allocation can be implemented using marginal taxes on labor income, bequests, and education expenditures that replicate the corresponding wedges. This approach provides a transparent link between the theoretical wedges derived in the planner's problem and empirically relevant policy instruments.

Education taxation and the human capital wedge. Using the first-order condition with respect to children's human capital and the envelope condition for the continuation value, the household's Euler equation for human capital investment in the decentralized economy can be written as

$$\begin{aligned} \frac{\partial U(c_t, l_t, h_{t+1})}{\partial c_t}\left(\frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}} + \frac{\partial T_t}{\partial h_{t+1}}\right) = &\frac{\partial U(c_t, l_t, h_{t+1})}{\partial h_{t+1}} \\ &+ \beta \mathbb{E}_t\left[\frac{\partial U(c_{t+1}, l_{t+1}, h_{t+2})}{\partial c_{t+1}}\left(\frac{\partial y_{t+1}}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}}\right)\right] \\ &- \beta \mathbb{E}_t\left[\frac{\partial U(c_{t+1})}{\partial c_{t+1}}\left(\frac{\partial T_{t+1}}{\partial y_{t+1}}\frac{\partial y_{t+1}}{\partial h_{t+1}} + \frac{\partial T_{t+1}}{\partial h_{t+1}}\right)\right]. \end{aligned} \tag{188}$$

Rearranging equation (188) yields a direct relationship between the marginal tax on human capital accumulation and the human capital wedge:

$$\frac{\partial T_t}{\partial h_{t+1}} = -\beta \mathbb{E}_t\left[\frac{\frac{\partial U(c_{t+1})}{\partial c_{t+1}}}{\frac{\partial U(c_t)}{\partial c_t}}\left(\frac{\partial T_{t+1}}{\partial y_{t+1}}\frac{\partial y_{t+1}}{\partial h_{t+1}} + \frac{\partial T_{t+1}}{\partial h_{t+1}}\right)\right] + \frac{\partial g(h_{t+1}, h_t)}{\partial h_{t+1}}\tau_{h,t}. \tag{189}$$

The human capital wedge $\tau_{h,t}$ is defined as

$$\tau_{h,t} \equiv \frac{\beta}{\partial g(h', h)/\partial h'}\mathbb{E}_t\left[\frac{u'(c_{t+1})}{u'(c_t)}\left(\frac{\partial y_{t+1}}{\partial h_{t+1}} - \frac{\partial g(h_{t+2}, h_{t+1})}{\partial h_{t+1}}\right)\right] + \frac{u'(h_{t+1})}{u'(c_t)\,\partial g(h_{t+1}, h_t)/\partial h_{t+1}} - 1. \tag{190}$$

Although the expression for the human capital wedge differs from that in Koeniger and Prat (2018) due to the presence of direct parental altruism toward children's education, the qualitative relationship between the marginal tax on human capital and the wedge remains unchanged. In particular, equation (189) highlights that, in a dynamic setting, a zero human capital wedge does not imply the absence of distortions in education investment, and a positive wedge does not necessarily translate into a positive marginal tax (Koeniger and Prat, 2018; Stantcheva, 2020).

This non-equivalence arises because human capital investment interacts with multiple overlapping distortions. First, education generates intergenerational spillovers by increasing the productivity of future generations. Second, education expenditures can serve as an implicit transfer mechanism, allowing dynasties to partially circumvent bequest taxation. Third, returns to human capital are indirectly taxed through the labor income tax schedule. The optimal education policy therefore depends on the joint configuration of labor, bequest, and education wedges rather than on the human capital wedge in isolation.

The first term on the right-hand side of equation (189) shows that the current marginal tax on human capital depends not only on future marginal tax rates on labor income and human capital, but also on how these taxes covary with the marginal utility of consumption. As emphasized by Bovenberg and Jacobs (2005), education subsidies and progressive labor income taxation should therefore be designed jointly and treated as "Siamese twins." Although their result is derived in a life-cycle framework, Koeniger and Prat

(2018) reach a closely related conclusion in a dynastic setting.

A key distinction in dynastic environments is that human capital investment affects not only contemporaneous earnings but also the productivity of future generations. As a consequence, the marginal tax on human capital incorporates both the marginal labor income tax and the marginal tax on next-generation human capital. Koeniger and Prat (2018) argue that, in order to avoid distortions, education expenditures should be fully tax-deductible. In particular, education subsidies are optimal when human capital investment raises the expected future tax burden and when the resulting changes in future taxation do not excessively increase consumption risk.

The introduction of parental altruism in the present framework does not overturn this policy recommendation. Under assumptions analogous to those in Bovenberg and Jacobs (2005)—namely, the absence of bequests, Cobb–Douglas production, and socially optimal human capital accumulation—equation (189) simplifies to

$$\begin{aligned}\frac{\partial T_t}{\partial h_{t+1}} = -\beta \mathbb{E}_t \left[\frac{\partial T_{t+1}}{\partial y_{t+1}} \frac{\partial y_{t+1}}{\partial h_{t+1}} - \frac{\partial T_{t+1}}{\partial h_{t+1}} \right] \mathbb{E}_t \left[\frac{u'(c_{t+1})}{u'(c_t)} \right] \\ - \beta \operatorname{Cov}_t \left(\frac{u'(c_{t+1})}{u'(c_t)}, \frac{\partial T_{t+1}}{\partial y_{t+1}} \frac{\partial y_{t+1}}{\partial h_{t+1}} - \frac{\partial T_{t+1}}{\partial h_{t+1}} \right). \qquad (191)\end{aligned}$$

A detailed derivation of equation (191) is provided in Appendix **??**. Stantcheva (2015) obtains analogous results in a life-cycle model and shows that the human capital wedge is shaped by the entire tax–transfer system. In particular, at the first-best optimum, education subsidies should equal the marginal labor income tax rate when the redistributive effect of education subsidies exactly offsets their efficiency cost, consistent with the mechanism identified by Bovenberg and Jacobs (2005).

The interaction of human capital taxation with labor supply and intergenerational transfers generates multiple, overlapping distortions. To address this complexity, **?** introduces the concept of the *net human capital subsidy*, under which education expenditures are fully tax-deductible, so that taxable income is reduced one-for-one by education spending. In a static setting, this corresponds to taxable income $y = wl - e$, where e denotes education expenditures.

In a dynamic framework, **?** shows that the optimal human capital wedge satisfies

$$\tau^*_{s,t} = \frac{\tau^*_{L,t}}{1 - \tau^*_{L,t}} \frac{\varepsilon^c_t}{1 + \varepsilon^u_t} \left(1 - \rho_{\theta s,t}\right), \qquad (192)$$

where $\tau^*_{L,t}$ is the labor wedge, ε^u_t and ε^c_t denote the uncompensated and compensated elasticities with respect to the net wage (holding savings fixed), and $\rho_{\theta s,t}$ is the Hicksian coefficient of complementarity between ability and human capital in the wage function. When $\rho_{\theta s,t} < 1$, the optimal human capital wedge and the labor wedge move in the same direction.

Overall, the central policy implication is that, under empirically plausible preferences and production technologies, education expenditures should be fully tax-deductible whenever individuals finance their own schooling.

In the following sections, we examine the quantitative implementation of two prominent education-financing instruments—income-contingent repayment loans and direct education subsidies—within the U.S. tax–transfer system.

H Income-Contingent Loans: Derivations

Household problem under income-contingent loans. The family's dynamic optimization problem is

$$\hat{W}(b,h,\theta) = \max_{c,l,h',b'} \left\{ U(c,l,h') + \beta \int_{\Theta} \hat{W}(b',h',\theta')\, dF(\theta'|\theta) \right\}, \tag{193}$$

subject to the intertemporal budget constraint

$$b' + c + g(h',h) = (1+r)b + \delta y^{1-t_y} + L_t(h',h) - D_t(b,y,h',L^{t-1},y^{t-1}). \tag{194}$$

At time t, parents may take an education loan to finance their children's human capital investment, while simultaneously repaying their own education debt through the income-contingent repayment function $D_t(\cdot)$. This structure mirrors student loan systems in Europe and Anglo-Saxon countries, where repayments depend on realized income—linked to human capital—and inherited wealth, proxied by bequests.

For transparency, the quantitative analysis focuses on the implementation for the first generation and abstracts from conditioning on pre-existing loan histories.

Figures 7 and **??** illustrate the implementation of income-contingent repayment loans under the current U.S. tax schedule, without requiring a comprehensive reform of the tax system. While the structure closely follows Koeniger and Prat (2018), the present framework incorporates parental altruism toward children's education, leading to qualitatively similar results with some differences in magnitude.

Figure 7 depicts how loan repayment varies with labor income for a representative dynasty with zero initial assets and highly educated parents. The upper panel shows the education loan amounts across income levels, while the lower panel presents the corresponding net repayment schedule. For incomes above $52,000, net repayments become mildly regressive. This pattern reflects the trade-off between insurance and incentive provision that characterizes the optimal design of education finance under a progressive labor income tax system (see equation (40)). The result is consistent with findings in Heathcote, Storesletten, and Violante (2017) and Koeniger and Prat (2018), who argue that optimal tax progressivity is lower than that observed in the current U.S. tax code.

Differences in net repayment patterns—especially with respect to bequests—arise from the inclusion of parental altruism. Wealthier parents not only leave larger bequests but also invest more heavily in their children's education, consistent with empirical evidence of overinvestment among high-income families. As a consequence, both low-income dynasties facing borrowing constraints and high-income dynasties with substantial bequests receive net education subsidies under the optimal policy.

To further simplify the analysis, we consider an alternative implementation in which loan repayments depend only on parental income, while loan amounts continue to depend on children's human capital according to

$$L_t(h',h) = g(h',h).$$

In this specification, the full cost of education is financed through student loans.

Figure **??** presents the resulting optimal repayment schedule. Families with labor income below $37,913 make no repayments, while repayments increase with income above this threshold. Although the education loan is formally taken out by parents, repayments are interpreted as being made by the child once active in the labor market. Accordingly, the optimal income-contingent repayment function faced by the working generation is given by

$$D_t(y) = \begin{cases} 0, & y \leq \$37{,}913, \\ -0.002964\, y^2 + 0.4455\, y - 12.63, & y > \$37{,}913. \end{cases} \tag{195}$$

H.1 Simplified Implementation of Income-Contingent Repayment Loans

We now consider a simplified implementation of income-contingent repayment loans (ICLs). Under this scheme, the family's optimization problem can be written as

$$\begin{aligned}\mathcal{L} = \max_{b',h',l} \Bigg\{ & U((1+r)b + L_t(h,h') - D_t(y,h') + y(l,h,\theta) - T_t^y(y) - b' - g(h',h),\, l,\, h') \\ & + \beta \int_{\Theta} \widehat{W}(b',h',\theta')\, dF(\theta'|\theta) \Bigg\} \\ & + \kappa\,(b' \geq \max\{-\phi g(h,h'), \underline{b'}\}) + \eta(h'-1),\end{aligned} \tag{196}$$

where $L_t(h,h')$ denotes the education loan taken by parents, $D_t(y,h')$ is the income-contingent repayment made by the next generation, and $T_t^y(y)$ is the labor income tax.

A detailed derivation of the optimal policies under income-contingent loans is provided in Appendix G. Here, we summarize the main quantitative implications for dynastic behavior and welfare. We compare the status quo—capturing the current U.S. education finance system—with an alternative economy in which education loans are repaid by the next generation, conditional on realized income.

Repayments follow the optimal net repayment schedule illustrated in Figure **??**. For middle-income households, repayments are increasing in labor income, while they are capped at the total loan amount and set to zero for the lowest- and highest-income families. In particular, exempting top earners from repayment encourages educational investment among high-ability dynasties and allows the planner to benefit from higher future labor income tax revenues. The repayment rule is given by

$$\tau_g(y) = -0.002964\, y^2 + 0.4455\, y - 12.63.$$

The quantitative effects of income-contingent loans are reported in Table 6. Average educational attainment increases by approximately 1.5 years, while the intergenerational correlation in human capital declines. These changes reflect improved access to education for individuals from low-income families and a reduction in borrowing constraints. Moreover, the human capital accumulation technology implies a lower marginal cost of education under ICLs.

Introducing income-contingent loans also raises aggregate welfare—measured as the sum of utilities across all dynasties—as well as total output, leading to higher tax revenues. Overall, the results in Table 6 indicate that income-contingent repayment loans improve social welfare and promote a more educated society.

Table 4 reports the effects of income-contingent repayment loans (ICLs) on educational inequality across earnings quantiles. The introduction of ICLs substantially reduces disparities in education costs between low- and high-income families: the average cost gap declines from \$75 under the status quo to \$59 with ICLs. This reduction reflects the insurance component of income-contingent repayment, which relaxes borrowing constraints and lowers the effective marginal cost of education for disadvantaged households.

At the same time, dispersion in educational attainment increases modestly. The schooling gap between

Table 6: Quantitative Effects of Income-Contingent Loans

Variable	**Data**	**Status Quo**	**(With ICL**
Average years of schooling S	12.86	12.52	14.06
Correlation (S', S)	0.46	0.47	0.45
Intergenerational earnings elasticity	0.45	0.44	0.45
Average net cost of an additional year of schooling	\$13,845	\$13,194	\$11,783
Bequest–wealth ratio	0.008	0.007	0.004
Education expenditures–income ratio	0.01–0.03	0.02	0.02
Aggregate welfare		1.8161×10^5	3.0590×10^5
Aggregate output	\$$14.99\times 10^{12}$	\$$1.2656\times 10^{12}$	\$$1.5256\times 10^{12}$
Aggregate income tax revenue (% of output)	10.67%	15.28%	18.33%

Notes: Model statistics are reported for $\vartheta = 0.0005$. Aggregate output in the data corresponds to U.S. GDP in 2010 (\$14.99 trillion, World Bank). The average individual income tax rate in the data equals 13.29 percent (Tax Foundation). Aggregate welfare is measured as the sum of utilities across all dynasties.

the upper and lower earnings quantiles rises from 0.15 to 0.17 years. This pattern reflects heterogeneous take-up and investment responses to ICLs: while access to education improves broadly, high-income and high-ability families—who face lower repayment risk and expect higher returns—expand educational investment by slightly more than low-income families. As a result, ICLs reduce financial inequality in education while modestly increasing inequality in years of schooling.

Overall, these findings highlight a trade-off inherent in income-contingent loan schemes: by lowering financial barriers and providing insurance, ICLs promote access and efficiency, but they may also amplify differences in educational attainment driven by underlying heterogeneity in ability and expected returns.

I Household Problem with Education Subsidies

This appendix provides the formal characterization of the household problem in the presence of education subsidies.

I.1 Problem Formulation

We consider a decentralized dynasty that chooses consumption c, labor supply l, savings b', and investment in children's human capital h'. The government provides an ad-valorem education subsidy τ_g, which reduces the private cost of human capital investment from $g(h', h)$ to $(1 - \tau_g)g(h', h)$.

The household's dynamic optimization problem is given by

$$\hat{W}(b, h, \theta) = \max_{b', h', l} \left\{ U(c, l, h') + \beta \int_{\Theta} \hat{W}(b', h', \theta')\, dF(\theta'|\theta) \right\}, \tag{197}$$

subject to the budget constraint

$$b' + c + (1 - \tau_g)g(h', h) = (1 + r)b + y(l, h, \theta) - T(b, y, \theta), \tag{198}$$

and the borrowing constraint

$$b' \geq -\phi(1 - \tau_g)g(h', h). \tag{199}$$

We assume that human capital investments are not directly taxed, i.e. $T(h') = 0$.

I.2 Lagrangian

The Lagrangian associated with the household problem is

$$\mathcal{L} = \max_{b',h',l} \Bigg\{ U\big((1+r)b - T^b(b) + y(l,h,\theta) - T^y(y) - b' - (1-\tau_g)g(h',h),\, l,\, h'\big) \\ + \beta \int_{\Theta} \widehat{W}(b',h',\theta')\, dF(\theta'|\theta) \\ + \kappa\,[b' + \phi(1-\tau_g)g(h',h)] + \eta(h'-1) \Bigg\}, \tag{200}$$

where κ and η are the multipliers on the borrowing constraint and the lower bound on human capital, respectively.

I.3 First-Order Conditions

The first-order condition with respect to savings b' is given by

$$\frac{\partial U}{\partial c_t} = \beta \mathbb{E}_t \left[\frac{\partial U}{\partial c_{t+1}} (1 + r - T^{b\prime}_{t+1}) \right] + \kappa_t. \tag{201}$$

The first-order condition for human capital investment is

$$(1-\tau_g)\left(\frac{\partial g(h_{t+1},h_t)}{\partial h_{t+1}} + \beta \mathbb{E}_t \left[\frac{\frac{\partial U_{t+1}}{\partial c_{t+1}}}{\frac{\partial U_t}{\partial c_t}} \frac{\partial g(h_{t+2},h_{t+1})}{\partial h_{t+1}} \right] \right) = \frac{\frac{\partial U_t}{\partial h_{t+1}}}{\frac{\partial U_t}{\partial c_t}} \\ + \beta \mathbb{E}_t \left[\frac{\frac{\partial U_{t+1}}{\partial c_{t+1}}}{\frac{\partial U_t}{\partial c_t}} \frac{\partial y_{t+1}}{\partial h_{t+1}} \left(1 - \frac{\partial T_{t+1}}{\partial y_{t+1}} \right) \right] + \kappa_t \phi (1-\tau_g) \frac{\partial g(h_{t+1},h_t)}{\partial h_{t+1}}. \tag{202}$$

The first-order condition for labor supply is

$$\frac{\partial U}{\partial c_t} \frac{\partial y_t}{\partial l_t} \left(1 - \frac{\partial T_t}{\partial y_t} \right) = \frac{\partial U}{\partial l_t}. \tag{203}$$

I.4 Discussion

The education subsidy τ_g reduces the marginal cost of human capital investment by scaling the cost function. In addition, it relaxes the borrowing constraint by lowering the amount of resources that must be financed privately. As a result, subsidies affect both the intensive margin of education investment and the ability of constrained households to finance education.

When the borrowing constraint binds, the multiplier κ_t enters the human capital condition, increasing the marginal value of relaxing the constraint. In this case, education subsidies have an additional effect through the borrowing channel.

In the main text, we focus on the unconstrained optimality condition and abstract from these additional terms for expositional clarity.